



Computer Science

AS & A Level | AQA | 7516 & 7517



2015 specification
first exams in 2017 (2016 for AS)

A LEVEL

AQA

Course Companion

for A Level AQA Computer Science

Includes AS and A Level

Update v1.1 – 29th September 2016

zigzageducation.co.uk

**POD
6061**

Publish your own work... Write to a brief...
Register at publishmenow.co.uk

Contents

Thank You for Choosing ZigZag Education	ii
Teacher Feedback Opportunity	iii
Terms and Conditions of Use	iv
Teacher's Introduction.....	v
Full Contents Listing.....	vi

1. Programming

2. Data Structures

3. Algorithms

4. Theory of Comp

5. Data Represent

6. Computer Syste

7. Computer Orga

8. Consequences o

9. Communication

10. Databases

11. Big Data

12. Functional Pro

13. Systematic Ap

Programming Cha

Answers and Solu

INSPECTION COPY

COPYRIGHT
PROTECTED



Teacher's Introduction

This resource covers all the theory needed for the A Level AQA Computer Science for first teaching in September 2015 – with the first exams in June 2017.

Each main topic of the specification is given its own section in the resource.

- | | |
|---|----------------------------|
| 1. Programming | 8. Consequences |
| 2. Data Structures | 9. Communication |
| 3. Algorithms | 10. Databases |
| 4. Theory of Computation | 11. Big Data |
| 5. Data Representation | 12. Functional Programming |
| 6. Computer Systems | 13. Systematic Design |
| 7. Computer Organisation and Architecture | |

Within each section there are student notes covering the specification content and descriptions of theory supported with examples, fact boxes, diagrams, images and

In addition, more generic pseudocode, code snippets are included for the following

- Visual Basic .NET
- C#
- Python
- Pascal/Delphi
- Haskell (*Functional Programming* topic only)

Questions and tasks are interspersed throughout the guide to test and develop understanding. There is also a separate set of high-level and assembly programming tasks and combine different programming concepts to test their skills as a whole.

Answers/solutions are included at the back of this resource to save the teacher a comprehensive set of definite answers. In some cases, there are equally valid answers have been given.

As this companion also includes all the content needed for the separate AS Computer Science (first teaching September 2015, with the first exams in June 2016), content which is *only* relevant to AS is indicated using the dotted border and **A LEVEL** stamp, as shown here.

This is designed to assist co-teaching between the levels.

Update v1.1 – September 2016

A number of improvements, including (but not limited to):

- High-level and pseudo code fixes
- Haskell code added to topic 12 (Functional Programming)

Free Updates!

Register your email address to receive any future free updates to this resource or other Computer Science resources you have purchased, and details of any promotions for you.

* resulting from minor specification changes, suggestions from teachers and peer reviews, or occasional errors reported by customers.

Go to zzed.uk/freeupdates

INSPECTION COPY

COPYRIGHT
PROTECTED



Full Contents Listing

1. Programming

- 1.1 Programming.....
- 1.2 Programming paradigms.....

2. Data Structures

- 2.1 Data structures and abstract data types.....
- 2.2 Queues.....
- 2.3 Stacks.....
- 2.4 Graphs.....
- 2.5 Trees.....
- 2.6 Hash tables.....
- 2.7 Binary trees.....
- 2.8 Vectors.....

3. Algorithms

- 3.1 Graph traversal.....
- 3.2 Tree traversal.....
- 3.3 Reverse Polish notation (RPN).....
- 3.4 Searching algorithms.....
- 3.5 Sorting algorithms.....
- 3.6 Dijkstra's shortest path algorithm.....

4. Theory of Computation

- 4.1 Abstraction and automation.....
- 4.2 Regular languages.....
- 4.3 Context-free languages.....
- 4.4 Classification of algorithms.....
- 4.5 A model of computation.....

5. Data Representation

- 5.1 Number systems.....
- 5.2 Number bases.....
- 5.3 Units of information.....
- 5.4 Binary number systems.....
- 5.5 Information processing systems.....
- 5.6 Representing images, sound and other data.....

6. Computer Systems

- 6.1 Hardware and software.....
- 6.2 Classification of programming languages.....
- 6.3 Types of program translator.....
- 6.4 Logic gates.....
- 6.5 Boolean algebra.....

INSPECTION COPY

COPYRIGHT
PROTECTED



7. Computer Organisation and Architecture

- 7.1 Internal hardware components
- 7.2 The stored program concept
- 7.3 Structure and role of the processor and its components
- 7.4 External hardware devices

8. Consequences of Uses of Computing

- 8.1 Individual, social, legal and cultural issues and opportunities

9. Communication and Networking

- 9.1 Communication
- 9.2 Networking
- 9.3 The Internet
- 9.4 Transmission Control Protocol / Internet Protocol

10. Databases

- 10.1 Conceptual data models and entity relationship modelling
- 10.2 Relational databases
- 10.3 Database design and normalisation techniques
- 10.4 Structured Query Language (SQL)
- 10.5 Client-server databases

11. Big Data

- 11. Big Data

12. Functional Programming

- 12.1 Functional programming paradigm
- 12.2 Writing functional programs
- 12.3 Lists in functional programming

13. Systematic Approach to Problem Solving

- 13.1 Aspects of software development

Programming Challenges

- Programming Challenges
- Assembly Programming Challenges

Answers

**COPYRIGHT
PROTECTED**



1. Programming

It is important to develop both practical skills and understanding of the theory behind them. This will be assessed in this course. Although there are many different languages – all with different syntax – they all share many of the same fundamental concepts, even if they do work in different ways.

This section covers:

1.1 Programming.....p1 1.2 Programming Par

1.1 PROGRAMMING

DATA TYPES

In order to run efficiently, computers need to be able to handle all forms of data.

When you declare a *variable* (see p.4), you must also declare a data type. This gives the computer an understanding of how much memory needs to be allocated as well as what operations can be applied to an item of data. For example, you cannot store an integer in a variable designated for storing text and vice versa.

Before you can start programming you must start with a blank canvas and start with the basics.

Language-defined data types

Integer	Any whole number (inclusive of negatives and zero). For efficiency, computers will offer a varying size. In increasing order these are 'short', 'int', 'long' when accuracy isn't of high priority.
Real	Often referred to as ' <i>float</i> ', this is any number within a <i>range</i> and can be user defined. These can be whole numbers and contain a <i>mantissa-exponent</i> form. As with integers there are two sizes: 'float' and 'double' used where accuracy is of high priority (e.g. when dealing with scientific data).
Boolean	Stores whether a condition is TRUE or FALSE. Default is set to FALSE. Some programming languages that do not support Boolean variables use 0 for FALSE and 1 for TRUE.
Character	This can contain any keyboard character, inclusive of special characters.
String	A set of characters: usually used to store a representation of text.
Date/Time	A representation of a moment in time. Can be used to return the current date and time.
Pointer	Sometimes referred to as ' <i>reference</i> '. Often used in linked lists. It stores the memory location of the next item in the list. Not all languages support the pointer construct.

**COPYRIGHT
PROTECTED**



User-defined data types

You can also declare your own type. These are called *user-defined* types. The *language-defined* types to make coding more efficient.

Enumerated

An *enumerated* data type is one that is in the form of a list. This could be a list of months in a year. Often these data types will be used for comparison; for instance, given a football player to see whether they are in the list of players in a team. It is straightforward when using an enumerated data structure. A problem with them is that the source code of the program and as a result cannot be changed once they are defined.

Sub-range

A *sub-range* data type defines a sub-range of elements from an enumerated data type. If a hierarchy of structure exists, for example, in the football team example you could define a sub-range for goalkeepers, defenders, midfielders and strikers.

Sets

A *set* is a structural data type and is the same as the mathematical idea of a set. It is a collection of elements, and an element can be a member of a set. For example, given the numbers from 1 to 100 you could define a set of odd numbers. Once this set had been defined you could use it to find the odd numbers.

```
Numbers1to100 = 1, 2... 100
evenNumbers = 2, 4... 100
oddNumbers = Numbers1to100 - evenNumbers
```

Arrays

An array is a data structure that can be used to hold elements of data of the same type. The elements can be retrieved later in the program's execution. The simplest type of array is a one-dimensional array. It has a given length, but no depth. You need to be able to search the array using an index. Arrays are *indexed* from the integer 0 to the user-defined length. An example of an array can be seen below.

It is also possible to have an array of multiple dimensions, most commonly two-dimensional. The first dimension is the number of directions you can index information from, so you can go *across* or *up and down*. For example, you could have a two-dimensional array of numbers. You can access them using their indexes to retrieve the information.

An example of use of a simple one-dimensional array containing integers would be to store a list of numbers and then sort them.

Pseudo	.Net	C#
Numbers[100] Numbers[56] ← 72	int[] Numbers = new int[100]; Numbers[56] = 72	int[] Numbers = new int[100]; Numbers[56] = 72

Interpretation of an array called *Numbers* – note index 56 has been assigned the value 72

0	1	2	3	4	...	55
55	102	11	87	65	...	61

COPYRIGHT
PROTECTED



Records

A record is a structural data type and is one that can provide a structure with which we were required to save the name, house number and postcode of 100 customers. This way each element of the array would contain a single set of details, i.e.:

```
Customer[72] =   Name: Fred Bloggs
                  Housenumber: 3
                  Postcode: BS10 5BY
```

Records or structured data types are defined by the user by building a declaration and then declaring a variable of that structure type. The program can then reference the variable using the notation.

Python	
<pre>Declare a structure as record name: string houseNum: integer postcode: string customers[100] as recStructure</pre>	<pre>Public Structure recStructure Public name As String Public houseNum As Integer Public postcode As String End Structure Dim Customers[100] as recStructure Customers[72].name = "Fred Bloggs" Customers[72].houseNum = 3 Customers[72].postcode = "BS10 5BY"</pre>
C#	
<pre>Struct recStructure { Public string name Public int housenum Public string postcode } recStructure[] customers = new recStructure[100]; Customers[72].name = "Fred Bloggs" Customers[72].houseNum = 3 Customers[72].postcode = "BS10 5BY"</pre>	<pre>type recStructure = struct name:string; housenum:integer; postcode:string; end; var customers: array<recStructure> = new array<recStructure>(100); Customers[72].name = "Fred Bloggs" Customers[72].houseNum = 3 Customers[72].postcode = "BS10 5BY"</pre>

Questions: Data Types

- What data types would best fit the following?
 - Welcome to England (1 mark)
 - 5 SE (1 mark)
 - 19 (1 mark)
 - 9001 (1 mark)
 - 4 (1 mark)
 - 17 Oct 1992, 3:44AM (1 mark)
- A bank is creating a new system that deals with the accounts of their customers. Identify suitable data types for the following variables and give a reason.
 - firstName (1 mark)
 - accountBalance (1 mark)
 - gender (1 mark)
 - hasOverdraft (1 mark)
 - dateOfBirth (1 mark)
 - sortCode (1 mark)

COPYRIGHT
PROTECTED



Variables and constants – explanation, declaration and use

Variables are parts of the program that allow it to change and perform *computations*. If a value of output from a program would never change, as the values would be fixed as constants. Fixed-value variables; where variables can change at run-time, constants are used in the source code. Both are *mnemonics* for a location in memory that need to be declared.

Language type and declaration

How you declare a variable (or constant) depends largely on what language you are using. It governs how the language performs 'type checking' – the act of preventing type errors. An action is performed on a construct, i.e. attempting to divide a string. In the case of declaring variables in two *strongly typed* languages. In dynamically typed languages, you do not declare a variable then use it. What actions can be performed on the variable are determined at compilation.



Variable		
VB.NET	Dim variableName As dataType	Const constantName As dataType
C#	dataType variableName;	Const dataType constantName;

Programming conventions and standards

As well as the rules of a programming language, there are also conventions and standards that you should follow throughout your programming career. While there are universal naming conventions (meaningful identifiers, no spaces, no special characters), there are also some language-specific conventions.

Indenting your code increases readability and allows the human eye to trace the flow of the program. It is a good practice to try to indent your code, even when writing pseudo language. *Python, use indentation to define how the program is run by defining blocks of code rather than using keywords.*

Naming conventions are another technique for improving readability of your code. There are several techniques; these are called 'camelCase' and 'PascalCase'. PascalCase is where the first letter of your identifier is upper-case while the rest of the word is lower-case; this is used for classes, subroutines and protocols. camelCase is where the first letter of the identifier is upper-case and the subsequent letters are lower-case; this is used for all other structures.

Did you know?!

Although these programming concepts are not enforced by any programming language, they are a good practice as even these simple ideas vastly improve readability. In fact, without these conventions, object-oriented programs can become very hard to read and understand.

COPYRIGHT
PROTECTED



Use of assignment

Assignment is one of the most fundamental operators as it allows you to change the value of a variable at any time. The way in which it is carried out is given the form *construct = expression* or *expression to construct*.

Pseudo	VB.NET
<pre>petrolCost ← 65.0 carHire ← 125.0 totalCost ← carHire + petrolCost</pre>	<pre>Dim petrolCost As Double = 65.0 Dim carHire As Double Dim totalCost As Double carHire = 125.0 totalCost = carHire + petrolCost</pre>
Python	Pascal
<pre>petrolCost = 65.0 carHire = 125.0 totalCost = carHire + petrolCost</pre>	<pre>var petrolCost: real; carHire: real; totalCost: real; totalCost:= petrolCost + carHire;</pre>

Note: in the VB.NET, Pascal and C# examples it would be equally accurate to assign the value to the variable after it is declared, as with petrolCost.

Iteration

It is quite common in programming to want to perform a certain task a fixed number of times or until a condition is met. Although it is possible to write out the code that many times, it is a bit cumbersome and rather impractical. To combat this, programmers can use the *loop* constructs.

FOR loops

The first loop you will learn about is the *FOR* loop. The FOR loop will run a set of statements a fixed number of times. The loop will use a variable to count the number of times it runs. For example, to check through a list of test scores to see how many students scored 90% the loop would look at each student's score, calculate their percentage and then output the result.

Pseudo	VB.NET
<pre>FOR i ← 1 TO 10 OUTPUT "i " ENDFOR</pre>	<pre>For i = 1 To 10 Console.Write(i & " ") Next End For</pre>
Python	C#
<pre>For i in xrange(1, 10) # using 'xrange' to generate a dynamic for loop Print i</pre>	<pre>For i:= 1 to 10 do Begin WriteLn(i); End;</pre>

It is worth noting in the above example that in C# the FOR loop will automatically increment the counter by 1 so that it doesn't require an 'End For'. In the Python code, 'xrange' is used to generate a dynamic for loop. In the Pascal code, 'writeln' is used to output the result dynamically so that, should the loop exit early, it uses much less memory.

WHILE loops

The WHILE loop is a very important basic structure. The syntax begins with a condition that will become the limit of the loop's run-cycle; once the condition is met the loop will run. If the condition is already met when the loop is called the code won't be run at all; it is only when the condition is not met that the code will run. An example of this can be seen when reading text from a file (see 2.1) where a WHILE loop is used to condition the reader to continue reading while the end of the file has not been reached.

COPYRIGHT
PROTECTED



Pseudo	VB.NET
<pre>Var ← 0 WHILE var <= 5 OUTPUT var Var ← var + 1 ENDWHILE</pre>	<pre>Dim i As Integer = 0 While I <= 5 Console.WriteLine(i) i = i + 1 End While</pre>
Python	
<pre>i = 0 # assign the count before it can be used While i <= 5 Print i i = i + 1 # if i has not been assigned i will not compile correctly.</pre>	<pre>Var i:integer; i:=0; While i<= 10 do Begin Writeln(i); i:=i+1; End;</pre>

In some instances you might be unsure as to whether you should use a FOR loop. Sometimes it is clear that the loop should execute a finite, known amount of times and a FOR loop should be used (*definite iteration*), but occasionally a loop might only need to execute until a certain condition has been met, such as a variable changing to a specific value, in which case a WHILE loop is more appropriate (*indefinite iteration*) – the only risk is that the condition will never be met and the loop will never exit the loop!

DO UNTIL loops

Also known as the *REPEAT UNTIL* loop, this iterative technique uses the same syntax as a WHILE loop except that the condition is evaluated at the end of the block. This means that the loop will always be executed at least once before it is evaluated against the control condition. If the condition is true; in this case it will continue to loop until the variable i has a value greater than 5.

Pseudo	VB.NET	C#
<pre>var ← 0 REPEAT OUTPUT var var ← var + 1 UNTIL var >= 5</pre>	<pre>Dim i As Integer Do Console.WriteLine(i) i = i + 1 Loop Until i >= 5</pre>	<pre>int i = 0; Do Console.WriteLine(i) i = i + 1; While (i >= 5)</pre>

Although there isn't a built-in *DO UNTIL* loop for Python, there is a way you can achieve the same result. It requires a little bit of ingenuity. By combining a *WHILE* loop with an *IF* statement and a *BREAK* command, you can create the iteration yourself.

Python	
Syntax	
<pre>while True: do_something() if condition(): break</pre>	<pre>while true: Print i i = i + 1 if i >= 5:</pre>

**COPYRIGHT
PROTECTED**



Breaking out of a loop

If you are using a FOR loop to iterate through a process and you are using a *break* command (for example) then it may also be useful to be able to 'break' the loop if the computer will continue to iterate through the loop until it reaches the end. The *break* command and it is placed after your condition variable is met. However, when using the *break* command to ensure the readability of code. Multiple breaks out of loops may be down to ask the question of whether a flag, better logic or a different kind of loop is

GoTo

It is possible to use *GO TO* to create loops and in other situations, such as *break* would do well to avoid using them. *GO TO* loops produce 'sloppy' and unreliable code wherever possible; in fact, the only time *GO TO* should be used is to cater for the error reporting code) or when *patching* code in post-release updates.

Selection

Selection is a control construct that is used as a control mechanism. A control statement set of values and determines the outcome. Examples of selection are the *IF* statement and the *CASE* select statement. A summary example can be found below.

IF selection

The common control flow statement is the *IF* statement. It is carried out by a condition. If the condition is TRUE then one portion of code is run. If the condition is FALSE

ELSE IF selection

Similarly to the *IF* statement, a condition is assessed for its value. However, the *ELSE IF* statement is carried out by the program. Look at the following example of the *break* command used to exit the loop once a condition has been met. The same can be achieved with a different kind of loop and altered conditions to test whether the number has reached a 'popty ping'.

Pseudo	
<pre> For i ← 1 to 15 a ← i MOD 2 b ← i MOD 3 If a AND b ← 0 Then # if mod division of both is zero OUTPUT "Popty Ping!" Exit FOR #exit after first found Else if a ← 0 Then # else if only a's mod division is zero OUTPUT "Pop!" Else if b ← 0 Then # else if only b's mod division is zero OUTPUT "Ping!" Else OUTPUT i # otherwise print i End IF End For </pre>	
VB.NET	
<pre> Dim a As Integer Dim b As Integer For i = 1 to 10 a = i mod 2 b = i mod 3 If (a = 0) And (b = 0) Then Console.WriteLine("Popty Ping!") End If Next i </pre>	<pre> For (int i = 1; i <= 15; i++) { int a = i % 2; int b = i % 3; if ((a == 0) && (b == 0)) { Console.WriteLine("Popty Ping!"); Break; } } </pre>

**COPYRIGHT
PROTECTED**



<pre>Exit For ElseIf a = 0 Then Console.WriteLine("Pop!") ElseIf b = 0 Then Console.WriteLine("Ping!") End If Next</pre>	<pre>else if (a == 0) { Console.WriteLine("Pop!") } else if (b == 0) { Console.WriteLine("Ping!") } else { Console.WriteLine("Neither") }</pre>
--	---

Python	
<pre>i = 0 for i in range(16): a = i % 2 b = i % 3 if a == 0 and b == 0: print('popty ping!') elif a == 0: print('ping!') elif b == 0: print('ping!') else: print('ping!')</pre>	<pre>Var a:integer; B:integer; begin for i:=1 to 15 do begin a:= i mod 2; b:= i mod 3; if (a=0) and (b=0) then begin writeln('popty ping!') break; end else if a = 0 then writeln('ping!') else if b = 0 then writeln('ping!') end; end end; end.</pre>

CASE selection

Sometimes there are multiple options to be considered, each one with a different letter out of the bag. If it is a *C* you will go to the cinema; if it is a *D* will go out for a run.

CASE selection has the option of an ELSE in the same way as IF selection. For example, if it is a *C* you will go to the cinema; if it is a *D* you will go out for dinner; otherwise you will stay in and watch television.

The following example shows how IF, ELSE and CASE statements are used.

<p>William is sitting at home and his mother says: 'Can you answer the telephone?'</p> <p>If it is Janice tell her I will call back later.'</p>	<pre>If caller = "Janice" then Message ("Mother will call back later.") End if</pre>
<p>William is sitting at home and his mother says: 'Can you answer the telephone?'</p> <p>If it is Janice pass me the phone, otherwise tell them I will phone back later.'</p>	<pre>If caller = "Janice" then Action ("Pass phone to Janice") else Message ("Mother will call back later.") End if</pre>
<p>William is sitting at home and his mother says: 'Can you answer the telephone?'</p> <p>If it is Janice pass me the phone, if it is Edith tell her I will be ready at 12, if it is Alfred tell him the time his cake is ready to collect, otherwise tell them I will phone back later.'</p>	<pre>Select Case caller Case caller = "Janice" Action ("Pass phone to Janice") Case caller = "Edith" Message ("Mother will be ready at 12") Case caller = "Alfred" Message ("Your cake is ready to collect") Case else Message ("Mother will call back later") End Select</pre>

**COPYRIGHT
PROTECTED**



Subroutines – procedures and functions

Subroutines are either *functions*, which return a value, or *procedures*, which return no value. Functions *must* be part of an expression but subroutines can also act as statements. Each is given an identifier and a list of parameters which are used to calculate the square-root calculator.

Pseudo	
<pre>"Enter an integer: " a ← INPUT OUTPUT ← squareroot (a)</pre>	<pre>Dim a As Integer Console.Write("Enter an integer: ") a = Console.ReadLine() Console.WriteLine(squareroot(a))</pre>
C#	
<pre>Console.Write ("Enter an integer: "); Int a = Console.ReadLine (); Console.WriteLine (math.sqrt(a)) ;</pre>	<pre>a = input('Enter an integer: '); # input is the key pressed on the keyboard Print Math.sqrt(a)</pre>
Pascal/Delphi	
<pre>Var a:integer; Begin Writeln('enter an integer'); Readln(a); Writeln(sqrt(a):0:3); End.</pre>	

Nested statements

Nested statements are when you have one set of statements *inside* another set. To elaborate further on the example above: pick a letter out of the bag. If it is a *C* you will go to the cinema (if it is raining you will drive, otherwise you will go for a run); if it is a *D* you will go out for dinner (if it is raining you will drive, otherwise you will go for a run); otherwise you will stay in and watch television.

Pseudo	
<pre>Select Case letter Case letter = "C" Action ("Go to cinema") Case letter = "D" If raining then Action ("Drive to restaurant") Else Action ("Walk to restaurant") End if Case letter = "R" Action ("Go for run") Case else Action ("Stay in and watch tv") End Select</pre>	<pre>Dim letter As Char Dim raining As Char Letter.ToUpper() Raining.ToUpper() Console.Write("Enter a letter: ") letter = Console.ReadLine().ToCharArray(0,1) Case letter = 'C' Console.WriteLine("Go to cinema") Case letter = 'D' Console.WriteLine("Pick a letter") raining = Console.ReadLine().ToCharArray(0,1) IF raining = 'Y' Then Console.WriteLine("Drive to restaurant") Else if raining = 'N' Then Console.WriteLine("Walk to restaurant") End If Case letter = 'R' Console.WriteLine("Go for run") Case else Console.WriteLine("Stay in and watch tv") End Case</pre>

**COPYRIGHT
PROTECTED**



C#	
<pre> Console.Write("Enter a letter: "); Char letter = Console.ReadKey(); Letter.ToUpper(); Switch (letter) { Case 'C': Console.WriteLine("Go to cinema"); Case 'D': Console.WriteLine("Is it raining? (Y/N) "); Char raining = Console.ReadKey(); Raining.ToUpper(); If (raining == 'Y') { Console.WriteLine("Drive to restaurant"); } Else if (raining == 'N') { Console.WriteLine("Walk to restaurant"); } Else Console.WriteLine("Incorrect input"); } Case 'R': Console.WriteLine("Go for run"); Case else: Console.WriteLine("Stay and watch tv"); } </pre>	<pre> Writeln('Enter a letter: '); readln(letter); case uppercase letter of 'C':writeln('Go to cinema'); 'D': begin writeln('Is it raining? (Y/N) '); readln(raining); If uppercase raining = 'Y' then WriteLn('Drive to restaurant'); Else if raining = 'N' then WriteLn('Walk to restaurant'); Else WriteLn('Incorrect input'); end; 'R': writeln('Go for run'); else writeln('Stay and watch tv'); end; </pre>

Note the use of the '.ToUpper' command. In the Unicode and ASCII character sets, there are upper-case and lower-case form. The command converts the input to upper case to make a CASE select. Python does not have a native 'switch-case' function built in, Pascal does.

Identifiers

Identifiers are the unique names given to elements such as variables and routines so that they can be identified. For this reason it is important that they are meaningful and relevant to the program, so that the program can be understood (potentially by other people than the original programmer).

Good use of identifiers is particularly important in complex programs which use a large number of variables and routines.

Questions: Programming Concepts

- Study the following pseudocode. It takes an array of results for a single test and stores it in the variable score. For each line (using the line number) of programming statement it is. (8 marks)


```

Procedure totalScore
a) Score = New Integer
b) Result = New Array
c) Pass = New Boolean
d) For i = 1 to 10
e)     Results = [3,7,5,7,3,6,8,4,2]
f)     PassBoundary = 30
g)     While (currentElement > maxElement)
h)         Score = Score + CurrentElement
        End While

```
- Complete the code by writing a nested CASE select in an IF statement to check whether the score is greater than the pass boundary; if TRUE begin the student to pass. The CASE select should calculate the student's grade by deducting the score and output their grade. The grade boundaries are A=30+, B=20+, C=10+, D=5+, E=0+.

COPYRIGHT
PROTECTED



ARITHMETIC OPERATIONS

Arithmetic operators

Arithmetic operators are basic functions you use when doing mathematics. The following table shows these arithmetic operators.

Operator	Meaning
=	Assignment
+	Addition
-	Subtraction
*	Multiplication
/	Division
^ or ↑ (common in C++)	Power

Note: you should never divide by zero, unless you want your program to give an error.

Modular arithmetic (MOD and DIV)

You know that 9 divided by 2 is **4.5**. You also know that 9 divided by 2 is **4** remainder **1**. You can work out these three numbers using /, MOD and DIV, i.e.:

```
a = 9 / 2;  
b = 9 DIV 2;  
c = 9 MOD 2;
```

When these three lines of code above are run, *a* is set to **4.5**, *b* is set to **4** and *c* is set to **1**. You use modular arithmetic in everyday life without even thinking about it; for example, if you work out how many days 50 hours is, you would work it out to be 2 days and 2 hours. If you work this out, it would look like this:

```
wholedays = 50 DIV 24;  
hoursleft = 50 MOD 24;
```

Brackets

Does $3 + 2 \times 5$ equal 13 or 25? Most languages follow BODMAS (brackets, or division, multiplication, addition, subtraction) as on a calculator which means it will do $3 + (2 \times 5) = 13$. You can use brackets anyway to ensure that the calculation is carried out in the order you want, but it's not slow down your program, but it makes it easier to understand the code.

Rounding

When you have floating point/decimal numbers you can round a value to a certain number of decimal places. The examples below illustrate the syntax of the round statement. For example, to round the value 3.14159 to four decimal places (~3.1416).

INSPECTION COPY

COPYRIGHT
PROTECTED



Pseudo	
<pre>Var ← 3.14159 Var2 ← round(var, 4) OUTPUT Var2</pre>	<pre>Dim myPi As Double Dim a As Integer a = Math.Round(pi, 4) Console.WriteLine(a)</pre>
C#	
<pre>Double myPi = 3.14159; Int a = Math.Round(pi, 4); Console.WriteLine(a);</pre>	<pre>myPi = 3.14159 a = round(pi, 4) print a</pre>
Pascal/Delphi	
<pre>const myPi:real = 3.14159; begin writeln(myPi:0:4); end.</pre>	<p><i>Note: in the Pascal version, the number of decimal places the value is rounded to is specified in the format string.</i></p>

Truncation

Like with rounding, truncation works on float/double numbers to remove value after the decimal point. It is used mainly in formatting when high precision is needed for operations but not for display. For example, the number pi is an irrational number which has no end digit. It has even been recited from memory to over 40,000 digits by a man in the UK. To use a value like pi, we have to truncate it to a manageable value without losing too much accuracy. In the following examples, we will see how to truncate a value to a specific number of decimal places.

Pseudo	
<pre>Var ← 3.14159 Var2 ← Truncate(var, 3) OUTPUT Var2</pre>	<pre>Dim myPi As Double Dim a As Integer a = Math.Truncate(pi, 3) Console.WriteLine(a)</pre>
C#	
<pre>Double myPi = 3.14159; Int a = Math.Truncate(pi, 3); Console.WriteLine(a);</pre>	<pre>myPi 3.14159 a = Math.trunc(pi, 3) print a</pre>
Pascal/Delphi	
<pre>const myPi:real = 3.14159; begin writeln(trunc(myPi*1000)/1000:0:3); end.</pre> <p>// There is no native truncate down command in Pascal so the above finds the number multiplied by 1000 to gain 3 decimal places truncated.</p>	<p><i>Note: 'myPi' has been used here because pi is a reserved keyword in other keywords of Pascal.</i></p>

Questions: Arithmetic Operations

- Answer the following:
 - 17 DIV 8 (1 mark)
 - 90 MOD 16 (1 mark)
 - ((16 DIV 2) * (6 MOD 4)) (1 mark)
 - 26 MOD 2 (1 mark)
- Write the pseudocode that can take two integer values and outputs whether they are divisible by each other without any remainder. (2 marks)

COPYRIGHT
PROTECTED



RELATIONAL OPERATIONS

Relational operations are the basis of making choices in mathematics. They are used to make the decision based on the situation. The following table contains the

Operator	Meaning
= or ==	Equal to
<> or !=	Not equal to
<	Less than
<=	Less than or equal to
>	Greater than
>=	Greater than or equal to

Questions: Relational Operations

- Answer the following (true or false):
 - $7 < 4$ (1 mark)
 - $4 > 1$ (1 mark)
 - $3.14159 \neq 3$ (1 mark)
- Write the pseudocode that can take two integer values and outputs whether they are divisible by each other without any remainder. (2 marks)

BOOLEAN OPERATIONS

In statements involving relational operations and conditions the following Boolean

Operator	Result
Expression AND Expression	AND only returns TRUE if <u>both</u> expressions are true.
Expression OR Expression	OR returns TRUE if either expression is true, and FALSE if neither is true.
NOT Expression	NOT returns the opposite of the expression, i.e. TRUE if it is false and FALSE if it is true.
Expression XOR Expression	XOR returns TRUE when the expressions are different and FALSE if they are the same.

Questions: Boolean Operations

- What would be the output for the following?
 - 'a XOR b' where *a* and *c* are *true* (1 mark)
 - 'c NOT d' where *c* and *d* are *false* (1 mark)
 - 'e AND g' where *e* is true and *g* is *false* (1 mark)

**COPYRIGHT
PROTECTED**



CONSTANTS AND VARIABLES

As stated earlier, all declarations are actually a shorthand representation of memory; this allocation is then given an identifier; this is the variable/constant. Variables and constants that require the variable name and data type to be declared are called *strongly typed*. If the computer tries to perform inoperable actions to the value.

Variables and constants are very similar. A constant is effectively the same as a variable, but it cannot be changed at run-time. This means variables can be assigned a value in the code and can be overwritten by a routine, but every time the program is restarted the original value is used. The value cannot be changed from the value in the source code.

An example can be seen every time you pay for something in a store or perform a calculation. The total price and the Value Added Tax (VAT). The total price for the transaction can be changed at run-time as more items are added to a transaction, whereas the VAT rate of the transaction cannot be changed by anyone, other than a manager or administrator. The VAT rate is a constant fixed rate. The pseudocode for the calculation is as follows:




```
const VATRate ← 1.2
```

```
... # run-time code for adding price of item to net total
Total ← varNet * constVATRate
```

Did you know?!

The use of mnemonics (naming variables) hasn't always been a feature of programming. In the early days, computers often required programmers to use the literal memory locations for computer memory addresses in RAM (before the introduction of offsetting!) every time they wanted to access a specific memory location.

STRING-HANDLING OPERATIONS



Strings are a series of characters. For example, a word or even an essay can be represented as a string. *Casting* is the process of converting between data types. Often you have information in one data type that you want to manipulate and will *cast* to another data type to perform an operation.

Some of the most common string handling operations are shown below:

Function	Description	Example
Length	Returns the length of <i>a</i>	Length("Computer")
Position(<i>a</i> , <i>b</i>)	Returns the position <i>a</i> in <i>b</i> , inclusive of special characters	Position("Computer", "m")
Substring(<i>a</i> , <i>b</i>)	Looks for string <i>a</i> within string <i>b</i> and returns TRUE if it is found	substring("Computer", "m")
Concatenate(<i>a</i> , <i>b</i>)	Joins string <i>b</i> on to the end of string <i>a</i>	Concatenate("Computer", "Science")

**COPYRIGHT
PROTECTED**



Alternatively to concatenation, you can also use string addition which function does the same as concatenation and can be seen below.

Pseudo	
exampleString ← 'Ex' + 'Ample' exampleString ← var1 + var2	exampleString = 'Ex' + 'Ample' exampleString = var1 + var2
C#	
exampleString = 'ex' + 'ample'; exampleString = var1 + var2;	exampleString = 'ex' + 'ample'; exampleString = var1 + var2;
Pascal/Delphi	
EXAMPLESTRING := 'EX' + 'AMPLE'; EXAMPLESTRING := VAR1 + VAR2;	

Character and character code conversions

To go from a character to a character code you can use the following:

Pseudo	
varAscChar ← 'a' varCharCode ← ConvertToAscCode(varAscChar)	Dim ascChar As Char Dim charCode As Integer
C#	
char ascChar = 'a'; int charCode = (int) ascChar;	ascChar = 'a' charCode = Ord(ascChar);
Pascal/Delphi	
ascChar:='a'; charcode:=ord(ascChar);	

To go from character code to the character representation you can use the following:

Pseudo	
varCharCode ← 97 varAscChar ← ConvertToChar(varCharCode)	Dim charCode As Integer Dim ascChar = chr(charCode)
C#	
int charCode = 97; char ascChar = (char) charCode	charCode = 97 ascChar = chr(charCode);
Pascal/Delphi	
charCode:=97; ascChar:= chr(charCode);	

Note: in pseudocode the required method of approaching this problem. If a character is given, you can use the ord function to get the character code. If a character code is given, you can use the chr function to get the character. Also, you are using a cast in the C# code to convert the integer to a character.

INSPECTION COPY

COPYRIGHT
PROTECTED



String conversion operations

You can also convert between data types; this is especially useful when reading data from a file.

Conversion	VB.NET	C#	Python
String to integer	CInt(string)	Convert.ToInt32(string)	Int (string)
String to float	Cdbl(string)	Convert.ToDouble(string)	Float (string)
Integer to string	CStr(integer)	Convert.ToString(integer)	Str (integer)
Float to string	CStr(double)	Convert.ToString(double)	Str (float)
Date/time to string	CStr(date/time)	Convert.ToString(Date/Time)	-
String to date/time	CDate(string)	Convert.ToDateTime(string)	-

Note: '--' has been used because DateTime is not a native data type in Python. If you use the 'datetime' class. Remember to import 'datetime' before you attempt to perform operations. In programming languages, the input is read as a string and must be converted before use.

Questions: String Handling

- What would be the results of the following built-in functions? (4 marks)
 - Length("Almost Impossible To Guess")
 - Round(656.3357, 2)
 - Length(Concat("Almost", "Impossible To Guess"))
 - Position('l', "Almost Impossible To Guess")
- Write code that asks the user for a number and prints the square of the number.

RANDOM NUMBER GENERATION

Another built-in function is the *random number generator*. Given a minimum and maximum value, it will generate a new random number when the function is called.

VB.NET	C#	Python
<pre>Dim newRand as New Random Dim x As Integer x = newRand.Next(1,10)</pre>	<pre>Random newRand = new Random(); Int x = newRand.Next(1,10);</pre>	<pre>From random import randint x = Randint(1,10)</pre>

In each case the syntax is very similar; note that in C# you call on the Random class.

Are the numbers actually random?

Often computers use a seed value to generate a sequence of what appears to be a random sequence. Programmers can set this seed value in many languages to be able to replicate the same sequence. This can be useful for testing applications. What you must remember about *randomness* is that it is hard for a logical computer to produce something that is truly random. This is because no matter how the random numbers are produced, the computer must rely on source code produced by a human to generate the numbers and it is impossible for humans not to introduce a portion of bias into a system. *Randomly generated* numbers are therefore given the title '*pseudo random numbers*'. These are numbers that appear random but have an underlying level of bias in how they've been produced.

INSPECTION COPY

COPYRIGHT
PROTECTED



EXCEPTION HANDLING

Exception handling can be used to help your program recover from run-time errors or divisions by zero. If an error occurs in the *Try-Catch* block it allows for the exception without the program crashing or losing any data.

Consider you'd like to check whether the user has input an integer:

Pseudocode	
<pre> "Enter an integer" Var ← Input Try #attempt conversion Var ← Convert to Integer Catch #output error "Input incorrect" End Try #end try </pre>	
<pre> Dim newInput As String Dim convertedInt As Integer Console.WriteLine("Enter an integer") newInput = Console.ReadLine() Try convertedInt = CInt(newInput) Catch (ex as Exception) Console.WriteLine("Input incorrect") End Try </pre>	<pre> Console.WriteLine("Enter an integer") String newInput = Console.ReadLine() Try { Int convertedInt = CInt(newInput) } Catch (Exception ex) { Console.WriteLine("Input incorrect") } </pre>
Python	
<pre> newInput = input('Enter an integer') try: int(newInput) except: # catches all errors print('Input incorrect') </pre>	<pre> println('Enter a number') readln(newInput); try numEntered:=StrToInt(newInput) except println('Input incorrect') end; </pre>

Learning to use these error handling methods now will save you a lot of stress and time when developing your own software during your project. The error handling methods you've been introduced to will help prevent almost any error that could occur in a system.

When your programming skills begin to develop beyond the basics you will use more advanced technologies; for example, you will be able to produce code that can interact with databases. Databases are renowned for being a minefield for amateur programmers, but with the correct knowledge and tools you can find it more comprehensible if something does go wrong.

Task: Error Handling

Write a program that asks the user for two integers that will be divided; build in error handling to protect against divisions by zero. If an exception occurs it prints the result as zero and continues.

INSPECTION COPY

COPYRIGHT
PROTECTED



SUBROUTINES

You have already looked at built-in functions such as *square root*, but now you will look at these structures. Subroutines are blocks of code which are independent of other code and can have their own variables, and they can be passed data using parameters to provide a return value. If a subroutine is called as part of an expression and returns a value it is called a *function*. If a *procedure* is a routine that is called as a statement which executes a section of code and returns an amount of results including none.

For example, you can use a procedure to open a file and a function to read the data from it. An example of a function can be seen below; this function returns the largest of two numbers.

Pseudo	
<pre> x = 17.0 y = 29.0 OUTPUT "The largest number is " & max(x,y) # the returned value of the function is added to the end of the output FUNCTION max (a, b) IF a > b THEN max = a ELSE max = b END IF Return Max END FUNCTION </pre>	<pre> Dim x As Double = 17.0 Dim y As Double = 29.0 Console.WriteLine("The largest number is " & max(x, y)) Function max (a As Double, b As Double) As Double IF a > b THEN max = a ELSE max = b End IF Return max End Function </pre>
C#	
<pre> Double x = 17.0; Double y = 29.0; Console.WriteLine("The largest number is " + max(x, y)); Static Real max (Real a, Real b) { Int Max; { IF (a > b) { Max = a; } ELSE { Max = b; } } Return Max; } </pre>	<pre> var x:real; var y:real; function max(a,b:real):real begin If a > b then max := a; Else max := b; end; begin x := 17.0; y := 29.0; writeln('The largest number is '); readln(max); writeln(max); end. </pre>
Python	
<pre> def max(a,b): max = 0 if a > b: max = a else: max = b return max x = 17.0 y = 29.0 print('The largest number is ' + str(max(x, y))) </pre>	<p>For an example of a function that reads data from a file, see p.21; for more on functions, see p.20.</p> <p>For more information on subroutines also see p.20.</p>

INSPECTION COPY

COPYRIGHT
PROTECTED



Did you know?!

To call a function or a procedure you simply type the identifier and pass any parameters. Notice in the VB.NET code that *a* and *b* are both defined as real numbers and that *real*; this has been done to show that you can cast within functions.

Look at the declaration for the function in the C# code. The function return type is *int*; this means that any variable you want to return must be returned as an integer. In this case, the function will still work properly.

In the Pascal/Delphi version, as soon as the function name is declared as a value it is returned. In simple code as above the multi-exit method is fine. In more complex routines it is better to use *exit* (see p.22) and then allocate at the end of the routine. Just like the 'break' command in C, which method to use is based upon ease of readability.

Questions and Exercises

- 1 Study the following code and describe in words how the function performs its task.

```
Function newSubroutine (Integer x)
    Answer ← x
    FOR var ← 1 to x
        Answer ← Answer * (x-var)
    END FOR
    RETURN Answer
```

- 2 What would be the output if the subroutine was passed the value '3'?

Procedures and functions as building blocks

Procedures and functions divide a program into building blocks. These basic blocks can be used to produce very complex programs and potentially reused in other projects. They make programs easier to read and more comprehensible, but also more space-efficient. Effective for any large-scale programming project.

Advantages of procedures and functions

There are a number of advantages of using procedures and functions:

1. Reduced amount of repeated code. For instance, if you know that during a program you will require to perform the operation $(a+b)*c$ very regularly then it makes sense to create a function that accepts the variables *a*, *b* and *c* as inputs and returns the value $(a+b)*c$.
2. Once a function has been written and is known to be correct, you know that it will work correctly for the rest of the program.
3. More than one programmer can work on the project, each on different parts of the program.
4. Once a function is finished the variables are deleted from memory, so that only the code needed to run the program is left.
5. Some quicker methods of sorting data use recursive functions.

INSPECTION COPY

COPYRIGHT
PROTECTED



PARAMETERS OF SUBROUTINES

Many functions require the calling program to pass information to them, which information, passed in the form of a variable, is called a *parameter* (or *argument*). The number of parameters. Parameters can also be a method of returning data. Parameters are placed inside the brackets after a function name:

Subroutine call	
<pre>myBase ← 6 myPower ← 2 OUTPUT (myExp(myBase, myPower))</pre>	<pre>Function myExp (tempPower) Counter ← tempPower Answer ← 1 While counter > 0 Answer ← ans Counter ← cou End While Return Answer</pre>

The words *myBase* and *myPower* above are the parameters for the function called *myExp*. It specifies the number of parameters (two, in this case), and most computer languages require parameters to be given (unless they are specified as optional).

Procedures and functions with interfaces

When creating large programs it is important to try to minimise the number of lines of code. Using modules and subroutines allows you to declare local variables within those code blocks. The single input and output interface ensures that the code starts and ends at the same place, making the code more intuitive. This restriction makes it much easier to debug than an unstructured approach. Consequently this method can be responsible for a single task.

RETURNING A VALUE FROM A SUBROUTINE

We've already explored how you saw that you can pass data to a subroutine to be used within that block. The data you pass is given a new declaration under a new temporary identifier that can be used within the block. So what if you want to access the data within a subroutine?

Some languages allow the transfer of variables by *reference* rather than by *value*. The function that worked out the maximum of two numbers used the result passed by the name of the function. However, it is also possible to use reference which will automatically alter the value in the routine that called it. This is useful for global variables or you may have to create a function to return it to the main program.

Consider a function that is used to calculate the area of a square. In the main program, assign the returned value to a variable; this must be of the same data type as the return value.

**COPYRIGHT
PROTECTED**



Pseudo	
<pre>"How long are the sides of your square?" a ← INPUT answer = squareCalc(a) FUNCTION squareCalc (sideLength) tempArea ← sideLength * sideLength Return var END FUNCTION</pre>	<pre>def squareCalc (sideLength): tempArea = 0 tempArea = sideLength * sideLength return tempArea a = input('How long are the sides of your square?') answer = str(squareCalc(a)) // note that Python uses str() to be considered a string</pre>
VB.NET (passing by Value)	VB.NET (passing by Reference)
<pre>Function squareCalc (sideLength As Double) As Double Dim tempArea As Double tempArea = sideLength * sideLength Return tempArea End Function Dim a As Double Dim answer As Double Console.WriteLine("How long are the sides of your square?") a = Convert.ToDouble(Console.ReadLine()) Answer = squareCalc(a)</pre>	<pre>Sub squareCalc (byRef sideLength As Double) Dim tempArea As Double tempArea = sideLength * sideLength Answer = tempArea End Sub Dim answer As Double Console.WriteLine("How long are the sides of your square?") answer = Convert.ToDouble(Console.ReadLine()) squareCalc(answer) // the variable a is passed by reference, so the variable a is changed</pre>
C# (passing by Value)	C# (passing by Reference)
<pre>Console.WriteLine("How long are the sides of your square?"); Double a = Convert.ToDouble(Console.ReadLine()); Double answer = squareCalc(a); Static Double squareCalc(Double sideLength) { Double tempArea; tempArea = sideLength * sideLength; Return tempArea; } // here a separate variable is used (answer) to store the result of using a passed as a value</pre>	<pre>Console.WriteLine("How long are the sides of your square?"); Double answer = 0; Double a = Convert.ToDouble(Console.ReadLine()); squareCalc(ref answer, a); Static Double squareCalc(ref Double answer, Double sideLength) { answer = sideLength * sideLength; } // the variable a is passed by reference, so the variable a is changed</pre>
Pascal/Delphi (passing by value)	Pascal/Delphi (passing by Reference)
<pre>var x:real; function squareCalc(a:real):real; begin squareCalc:= a*a; end; begin write('How long are the sides of your square?'); readln(x); writeln('The area is ', squareCalc(x):0:2); readln; end. // in the above function the variable x remains as the length and the function returns the area</pre>	<pre>var x:real; Procedure squareCalc(var a:real); begin a:= a*a; end; begin write('How long are the sides of your square?'); readln(x); squareCalc(x); writeln('The area is ', x:0:2); readln; end. // in this routine the variable x is passed by reference, so the variable x is changed through the procedure automatically when it changes the value of x</pre>

**COPYRIGHT
PROTECTED**



LOCAL VARIABLES

When you begin to use subroutines you'll discover it becomes more efficient increases *modularity* and decreases the volume of memory used. Modularity has been decomposed into individual problems; the aim is to have a subroutine process in a solution. Local variables cannot be called by anything outside of need to be passed as a parameter to the subroutine and any value from a local the main body. These variables only exist while the subroutine is being executed the main body of code and all memory is reallocated.

There is also the logistics of programming to consider. For example, on large a single team of programmers to complete the entire task so it is broken down on specific sections. This is to stop different teams from using the same identifiers otherwise cause a problem once the modules are combined.

GLOBAL VARIABLES

When you start learning to program in a language, most variables will be These are variables that can be called and are operable by all blocks of code simple programs this may seem easier, but consider the following:

- Global variables are assigned memory at run-time; this memory is open closes. If you're building a very large and complex program with many assigning a lot of system resources to variables that the user may open may choose to create a new item instead of reading from a file – memory from a file are made redundant. The end result is your program running
- As these variables are callable from all blocks they may be called and variable names are similar.

Questions: Procedures, Functions and Variables

- 1 What is the main difference between a procedure and a function? (1 mark)
- 2 What happens if you pass variable values into a subroutine in a different
3. In terms of memory and modularity, why is it considered bad practice to

INSPECTION COPY

COPYRIGHT
PROTECTED



ROLE OF STACK FRAMES IN SUBROUTINE CALLS

Whenever a routine is called, the computer allocates memory in a specialised 'stack'. When you use a function you'll allocate the returned value to a variable. The returned value is stored to the stack frame. This is the area where all the routine's parameters are stored until the routine ends, at which point the value is stored to the return variable. The value is then passed back to the main function and the stack frame is removed from the stack.

Stacks are a very useful data structure and they have many applications in computing, including calling procedures in programs. For example if a procedure `main()` calls a procedure `getchar()` then `main()` cannot continue until `getstring()` has finished. `getchar()` has finished; i.e.:



getChar()
getString()
main()
empty

Every time a procedure is called by another procedure it is pushed onto the stack. The stack grows downwards in memory, which the procedures need to be executed in an efficient way. If you have a procedure that calls itself indefinitely it will run out of stack space – try it!

RECURSIVE TECHNIQUES

Recursion is the ability that a subroutine has to call on itself to complete its task. Recursive solutions can be harder to produce but can often lead to very elegant solutions. A recursive solution has two parts: the *recursive* and the *limiter*. The recursive part is what calls itself another iteration and passes new variable values, whereas the limiter is what stops the infinite loop.

Recursive methods act as a loop that calls on itself and runs every line of code in the current call, and will pass the current result into the next call until the base case is reached.

One of the best examples of recursion is factorials. Factorials are given by the formula:

$$n! = n \times (n-1)!$$

This means a factorial is the product of a number times the factorial of the previous number.

$$4! = 4 \times 3! = 3 \times 2! = 2 \times 1!$$

$$4! = 4 \times 3 \times 2 \times 1 = 24$$

Answer: Calculate 4!

```
FUNCTION factorial (var)
IF var < 1 THEN # 1! = 1
    Return 1
Else
    Factorial var * (var - 1)
    # recall the function with new values
END IF
END FUNCTION
```

```
Dim answer As Integer
Dim base As Integer

Answer = factorial(4)

Function factorial(n As Integer) As Integer
    If n <= 1 Then
        Return 1
    Else
        Return n * factorial(n - 1)
    End If
End Function
```

**COPYRIGHT
PROTECTED**



C#	
<pre> Int base = 4; Int answer = factorial (base); Static Int factorial (Int n) { If (n <= 1) { Return 1 } Else { Return n * factorial (n - 1) } } </pre>	<pre> Base = 4 Answer factorial Def factorial (n) If n == 0: Return 1 Else: Return n </pre>
Pascal/Delphi	
<pre> function factorial(base:integer):integer; begin if base<1 then factorial:= 1 else factorial:=base * factorial(base-1); end; begin writeln(factorial(4)); end. </pre>	<p>In this example the the ELSE statement The limiter is the IF variable is less than producing an infinite</p>

Task: Recursive Techniques

The Fibonacci sequence is a set of numbers derived from the rule: $F_n = F_{n-1} + F_{n-2}$. Each number in the sequence is the sum of the previous two numbers of the sequence. This is a pattern of numbers that is found frequently throughout nature and even has applications in the design of computer components. For example, 9 times out of 10 the number of petals found on a newly blossomed flower will be a Fibonacci number, and in science it is said that any number can be written as the sum of unique Fibonacci numbers.

Using recursion, create a function that produces a list containing a Fibonacci sequence.

Questions: Recursive Techniques

Study the following pseudocode and answer the questions below.

```

FUNCTION MyFunction (Sum)
OUTPUT "enter an integer value: "
i ← READ VALUE
WHILE True
    IF i * 0 Then
        sum ← sum + i
        IF sum > 100 Then
            RETURN sum
        ELSE
            i ← i + 1
        END IF
    ELSE
        MyFunction(sum)
    END IF
END IF
END WHILE

```

- 1 What task does this function carry out? (1 mark)
- 2 Does this function enter an infinite loop? Explain your answer. (2 marks)
- 3 Describe the use of the stack in the above code. (1 mark)

**COPYRIGHT
PROTECTED**



1.2 PROGRAMMING PARADIGMS

STRUCTURED PROGRAMMING

In the early days of computing, computers were programmed by writing machine code. The advantage of allowing the programmer to directly control the computer by direct addresses and computer operations. To aid programmability the instructions were written in a language rather than binary. However, machine code is very difficult to read and so even experienced programmers would struggle to understand a large piece of code without a significant amount of time.

This led to the development of assembly code, which replaced the hexadecimal with mnemonics which were easier to read and understand. Most assemblers use labels (or variables) to signify memory addresses which would be referenced to a memory location. An assembly language, however, still operated at a very low level and became more complex as programs became larger.

One of the big problems with assembly language is the use of *GO TO* statements which are very difficult for people to follow. Structured programming developed as a way of providing a higher level of abstraction away from the operations of the computer to a level that is more understandable for humans to understand.

The structured programming paradigm encompasses procedural programming, object-oriented programming, and implies the ability to use structures such as IF statements and loops.

PROCEDURAL-ORIENTED PROGRAMMING

Procedural programming is a step forward in programming and provides a structured way of writing a program. Programs written in procedural languages are executed line by line in the order they are designed with a top-down view. From the top-down view, the program is seen as a series of steps written with procedures where each procedure performs a specific task. Program variables are used to store data which is local to each procedure. Bad programming paradigms, such as using *GO TO* statements, become unnecessary and are replaced by loops and procedures. Statements are grouped together and form procedures which perform a specific task. Having structure also automatically introduces recursion.

Many solutions can be broken down into a series of operations which can be performed in a logical order. However, more complex tasks and data structures lend themselves to more advanced programming methods. Often modern-day programming languages allow the programmer to use an object-based methodology.

OBJECT-ORIENTED PROGRAMMING

As procedural programs became more widespread, people started to notice that the procedures/functions associated with a task tended to be grouped together; this is the basis of object-oriented programming (OOP). At their most basic level, object-oriented programming concerns itself with the data for the *objects* you are trying to manipulate rather than the operations themselves. Before you start to use the object-oriented approach there are several key concepts to understand.

Classes

A class is an object definition. For example, a game might have a class 'goblin' which defines the goblin, such as name, health, weapon and colour and also the actions the goblin can perform, such as attack, defend, etc. Creating an object from a class invokes the constructor for that class.

1. Allocates and initialises the necessary memory
2. Assigns a label to that memory
3. Assigns values to (initialises) various properties as required (e.g. the goblin's name, health, etc.)

INSPECTION COPY

COPYRIGHT
PROTECTED



Objects

An *object* is an *instantiation* (or an *instance*) of a class. Each object will have the local to that object. For example, a computer can have states (on, off) and behaviour. Objects are created using a *constructor* and a *reference* that has been assigned to it.

Encapsulation

Encapsulation is where attributes and methods are 'wrapped' together into objects. The attributes and methods are kept together but their implementation details are *hidden* from one another. The communication of processes from other objects and classes is achieved by using the keyword *public*.

Encapsulation is applied:

1. If the internal complexity is not needed by other objects and doesn't change.
2. If you need to prevent changing an object from external objects.

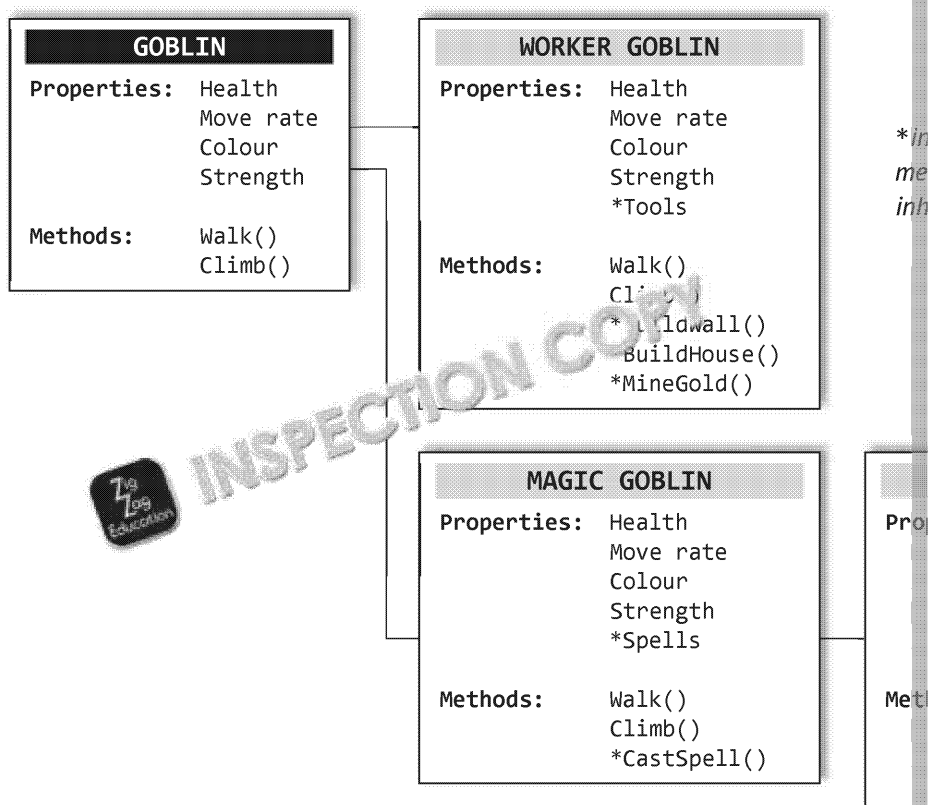
Inheritance

Inheritance is a relationship among classes wherein one class shares the structure and behaviour of another. This is similar to how children *inherit* the attributes of their parents.

Single Inheritance

Once a behaviour or characteristic is defined, all the categories beneath that category inherit that behaviour or characteristic. For example, the class *goblin* might define a goblin as having health, movement rate, etc. with methods for walking and climbing. You could then have *goblin* or *worker goblin*, where both then have inherited the properties (health, movement rate, etc.) and methods (walking and climbing) from their *parent* class (*goblin*), however each subclass can have its own properties, for example these two subclasses could have methods for casting spells or building houses respectively.

Extending this idea further, subclasses can be extended with further subclasses. For example, a *worker goblin* can be defined, with extra methods and properties on top of those defined in the *goblin* class.




COPYRIGHT
PROTECTED



Multiple Inheritance


Multiple inheritance occurs when a class inherits from more than one parent object-oriented language because many base classes can be set up from which to inherit. The following example of single inheritance is a very simple one. If an object inherits from a single inheritance then this limits the classes we could have; for example we could have: memory, output devices, input devices and processing devices.

Hardware – Single Inheritance



However, we have a problem if we want to bring in a new class called, for example, *Portable HDD* will acquire properties and methods from both the classes.

Hardware – Multiple Inheritance



When applied to object-oriented programming, the process is about building up a system together with the methods that accompany the data structures.

Aggregation

Aggregation in its everyday sense is similar to how programmers use it; it defines a relationship. In programming, there are two kinds:

1. *Association* is how objects are related without there being an owner. For example, a student can be taught by a single tutor and a tutor could teach many students. An association is not a strong relationship because if either is removed/deleted the other remains; just as a student moves on from a school doesn't mean the teacher is fired.
2. *Composition* is where the whole is defined by the relationship between the parts. For example, a house is the parent object/class then all children objects are removed. Consider a house is the parent class because it is the container in which all rooms are defined. If you remove a house it means you've destroyed the rooms, but you can change the rooms.

Polymorphism

Polymorphism refers to a programming language's ability to process objects of different type or class. This means that the code itself must be able to redefine methods for derived objects. For example, polymorphism would allow a programmer to define methods for any number of derived shape classes.

Method overriding

Method overriding is when you change the base characteristics of a class with derived code; this acts as an extension to the class. Take a look at the following example.

OOP: An example

The following code is written in pseudocode but the concepts are still the same. In this example you will be shown two classes that will be used together (*aggregation*). You will generate a version of an object-oriented design to create an object of a class and how you can use *constructors* to handle any variations in inputs.

Person.class	Person
<pre> CLASS STRUCTURE person private age private firstName private surname # Default Constructor STRUCTURE Person () age ← 0 firstName ← "No name set" surname ← "" END STRUCTURE # Partial Instantiated Constructor STRUCTURE Person (initialAge, personName) age ← initialAge firstName ← personName END STRUCTURE #Fully Instantiated Constructor STRUCTURE Person (initialAge, personName, personSurname) age ← initialAge firstName ← personName surname ← personSurname END STRUCTURE PROCEDURE GrowOlder () age++ END PROCEDURE END CLASS STRUCTURE </pre>	<pre> CLASS STRUCTURE Person #Declare new Person objects Person somePerson Person someOtherPerson Person someStrangePerson # Using the default constructor somePerson ← new Person() # using the partial constructor someOtherPerson ← new Person(25, "John") # using the fully instantiated constructor someStrangePerson ← new Person(30, "John", "Smith") #Calling a method on someStrangePerson someStrangePerson.GrowOlder() END CLASS STRUCTURE </pre>

You can use *UML* class diagrams without having to look at all the code behind them. *UML* diagrams convey all the information needed to understand how the system will form the final solution. They are a structural modelling technique used during the system life cycle (see Section 13).

**COPYRIGHT
PROTECTED**




Below is the view of a single entity.

Employee	← Class name
firstName(); lastName(); dateOfBirth(); salary(); holidayRemaining(); ...	← Attributes
GetHoliday(); SetHoliday(); SetSalary(); ...	← Methods

If you imagine that the class stores the information, it would be fine; however, and so must encompass not all members of staff, many would be on annual leave or cafeteria workers.

You could add another attribute but then not all members would have an hourly rate. The solution is to use 'SetSalary' to include

The full diagram might look something like the one below. Note the use of composition for aggregation. For example, *works_on* is composition.



Advantages and disadvantages of object-oriented design

Advantages	Disadvantages
<ul style="list-style-type: none"> Improved software maintainability Improved software stability Lower cost of development Higher-quality software Class code is reusable 	<ul style="list-style-type: none"> Harder to produce efficient code Approach is not suited for a small program size Slower program execution time Can be difficult to apply over a large program

Task: Object-oriented Programming

- Using the PersonGenerator class example, produce similar code that contains the first name, surname and bank balance of three bank accounts.
 - Improve on your code and write a subroutine that deposits £10 into an account.
- Hint:** Remember you can't have negative deposits.

**COPYRIGHT
PROTECTED**



2. Data Structures

Data structures play an important role in computer science. A good understanding of these particular, their relationship with certain programming techniques, is key to be able to solve problems. This section explores some structures available and how they are used.

This section covers:

2.1 Data structures and abstract data typesp1	2.5 Trees.....p11
2.2 Queues.....p11	2.6 Hash tables.....p14
2.3 Stacks.....p14	2.7 Dictionaries.....p17
2.4 Graphs.....p17	2.8 Vectors.....p17

2.1 DATA STRUCTURES AND ABSTRACT DATA TYPES

DATA STRUCTURES

All data types, from the most basic integer to the most complex tree, can be categorised into data types. Each category uses memory in a different way, and compilers must manage the memory accordingly.

- **Strong types** are the standard types, such as integer and character. They have a fixed amount of memory, and a fixed amount of memory is defined.
- **Static types** are those which require a fixed amount of memory, such as arrays. They are not included in the pre-defined types. For example, an array in C# is a static type.
- **Dynamic types** are those that may be expanded given the limitation of memory (in the case of files); for example, files and pointers (thus linked lists, etc.).

SINGLE- AND MULTI-DIMENSIONAL ARRAYS

The array is one of the most useful and fundamental data structures there are. It is a *matrix* of a single data type; a matrix is where you can store data into elements and retrieve the data from the element by using a unique identifier. Ensure you have a good understanding of arrays, as you will find that you can simplify many aspects of your programs using arrays, as you will find that you can simplify many aspects of your implementation (and it will help you if an arrays question comes up in the exam).

One-dimensional arrays

Storing data in several different variables becomes tedious and impractical with large amounts of similar data. The solution to this is an array. An *array* is a set of data stored consecutively in memory. If you wanted a set of 10 integers called X you might write:

Pseudo	VB.NET	Python
X [10]	int x [10]	x = [] # This creates a list of size 0

Note: in pseudocode you reference each term by X[0], X[1], X[2], ... X[9]. The number is known as array subscripts. Some languages would define the array as going from 0 to n-1, others define it as '0 to n' (e.g. pseudocode, Pascal/Delphi). Although the latter is a common method often makes array manipulation much simpler and is closer to the actual memory layout.

Suppose you stored the marks of 10 tests in an array. A common way to visualise this is:

81	75	90	64	68	72	69
X[0]	X[1]	X[2]	X[3]	X[4]	X[5]	X[6]

COPYRIGHT
PROTECTED



Most programming languages use loops to iterate through a data construct and you can access the memory location of a result directly if you knew what result. However, in most cases you will need to iterate through an array to find the result. Look at the following example code which iterates through the array x to find the result.

Pseudo	
<pre> OUTPUT "student with top marks: " For i 1 to 10 If x[i] > topScore then topScore = x[i]#1 end if next OUTPUT topScore #1 sets topScore to array index [i]</pre>	<pre> Dim topScore As Integer Dim i As Integer For I = 0 to 9 If x[i] > topScore then topScore = x[i] End If Next Console.WriteLine(topScore)</pre>
C#	
<pre> Int topScore = 0; For (Int I = 0; i = 9; i++) { If (x[i] > topScore) { topScore = x[i]; } } Console.WriteLine(topScore);</pre>	<pre> topScore = 0 For i in range(len(x)) if x[i] > topScore: topScore = x[i] print(topScore)</pre>
Pascal/Delphi	
<pre> var x:array [1..10] of integer; i:integer; topscore:integer; begin for i:= 1 to 10 do begin if x[i]>topscore then topscore:=x[i]; end; writeln(topscore); end.</pre>	<p><i>Note: Python does not use arrays and instead uses a list. Lists can hold multiple data types; they are flexible in the actions that can be performed on the data they can contain. This makes them slightly harder to use than arrays.</i></p>

Two-dimensional arrays

In most languages arrays can have more than one dimension, as many as 32. A dimension is simply a direction in which you can vary the specification of elements. With two-dimensional arrays you can vary along the columns and across the rows.

Declaration and use

The declaration is almost the same except we define the length of the array.

Pseudo	VB.NET	C#	Python
X [3, 3]	int x [3, 3]	Int x[,] = new int [3, 3];	x = [[]]*3

You can then assign values as you would with a single array. This is often used for the same data type that can be compared to each other. *Note: Python does not have arrays although a 2D array can be emulated through a list of lists.*

**COPYRIGHT
PROTECTED**



For example, in this table you can see the distances between varying cities. *context and do not appear in the array.* See if you can replicate the array below

	Plymouth	London	Edinbu
Plymouth	0	237	487
London	237	0	413
Edinburgh	487	413	0

Three-dimensional arrays

Take what you already know about arrays and think about how you would manage data that would be able to index data across the rows, down the columns and have a third dimension. This can be hard to show on paper, but it is a natural way of storing data that can be conveyed. For example, you could store the coordinates of a vector in a single-dimensional array, but you would be limited to having no control over direction. You could store the vector coordinates in a two-dimensional array, but you would still have no control over depth.

Three-dimensional arrays allow you to store real special dimensional data, and are used in many software and modelling software that relies heavily on vectors instead of bits.

Questions: Arrays

- 1 Consider the following single-dimension array and answer the questions below.

Array RawMarks

(0)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
31	18	27	32	9	28	11	17	21

- a) What value is retrieved by 'RawMarks[6]'? (1 mark)
 b) What is the output of the following algorithm? (2 marks)

```

FOR i ← 0 To Length(RawMarks)
  n ← 0
  IF RawMarks[i] > n Then
    n ← RawMarks[i]
  END IF
  OUTPUT n
END FOR
  
```

- 2 Study the following two-dimensional array about the distances between planets in the Solar System and answer the questions below. Planet names have been abbreviated.

Array PlanetaryDistance	0 (Mercury)	1 (Venus)	2 (Earth)
0 (Mercury)	0	0.33	0.61
1 (Venus)	0.33	0	0.28
2 (Earth)	0.61	0.28	0
3 (Mars)	1.13	0.80	0

- a) How would you access the values for the distance between Mercury and Venus?
 b) How would you access the values for the distance between Venus and Earth?
 c) Notice that the indexes PlanetaryDistance[2, 3] and [3, 3] are the same. How would you change their values so that they are correct? (1 mark)

INSPECTION COPY

COPYRIGHT
PROTECTED



FIELDS, RECORDS AND FILES

Very few programs do not utilise files at some part of their execution. Without effectively starting from the beginning and would require the data input even programs that are run as new each time tend to have user settings or preferences. Therefore the use and understanding of files is a vital part of programming. Text/ASCII files and binary files.

Text/ASCII files store the information in text or ASCII character format. If the another program (such as Notepad or Excel) then the file is stored as an ASCII file contains coded data which, without the context of the program, would not a text file might contain the data 'AB123' which has no meaning unless applied. The advantage of storing information in text files is that they can be created, altered the program. This can be useful in the development of a system to allow the

Binary files are stored as binary (often in hexadecimal format). These are but are directly compatible with the computer. Binary files tend to be more specific definitions which are difficult to convert the data into meaningful information.

Here is a text/ASCII file of CSV (comma-separated variable) data called *students.csv*. The data is in the format *surname, ID, course, college*.

Note that you call each line of the file a *record*, and you call each part of the record a *field*.

Before reading the file to be used in a program, you may wish to create a data structure to contain the data that is shown below. Once the structure is created you can then apply the I/O from the file to the structure to contain the data.

Pseudo	
<pre> NEW STRUCTURE ← Record # declare the structure # Declare the variable names and types Surname ← string ID ← integer Course ← Integer College ← String END STRUCTURE </pre>	<pre> Structure Record Public Public Public Public End Structure </pre>
C#	
<pre> struct Record { public string surname; public int ID; public int Course; public string College; } </pre>	<pre> Class recordStruct def __init__(self, self.surname; self.id; self.course; self.college; </pre>
Java/Delphi	
<pre> Type TStudent = Record Surname: string; ID : Integer; Course : Integer; College:string; End; </pre>	

INSPECTION COPY

COPYRIGHT
PROTECTED



Reading from a text file

Using input and output is very different depending on the programming language to be able to recognise and operate this function as it widely broadens what

For the C-family, in order to read from a text file you need to invoke the help *StreamReader*. Stream readers are a way of interfacing between the contents building, and use a string variable to read the text file line by line, whereas as Python have a class called *open* that defines what is being used and what

When reading from a file you must remember that all values returned from a includes numbers. In order to use the values being read you may need to cast

Pseudo	
<pre> CONST ← 'C:\MyDocuments\ReadingTest\file.txt' Line ← Length of one line filename ← file name USING StreamReader ← streamreader WHILE NOT EOF(filename) DO Readline(filename, OneLine) OUTPUT (OneLine) End while Close(myFile) </pre>	<pre> Dim FileRead Const filename Dim oneLine filename = "C:\MyDocuments\ FileReader Do Until File on Co Loop FileReader. Console.Re </pre>
C#	
<pre> string lineFromFile; streamReader reader; string fileName = @"C:\MyDocuments\ReadingFileTest.txt" // '@' is added as an escape char. reader = new StreamReader(fileName); While (!reader.EndOfStream) { lineFromFile = reader.ReadLine(); Console.WriteLine(lineFromFile); } Reader.Close(); </pre>	<pre> fileName = 'C:\MyDocuments\ file = open for line in print(l file.close(</pre>
Pascal/Delphi	
<pre> var fileIn:textfile; lineFromFile:String; begin assignfile(fileIn,'C:\MyDocuments\ReadingFileTest.txt'); reset(fileIn); // sets the program ready to read from fileIn and moves the pointer to the top while not eof(filein) do begin readln(fileIn,lineFromFile); writeln(lineFromFile); end; closefile(fileIn); end. </pre>	

INSPECTION COPY

COPYRIGHT
PROTECTED



Writing to a file

Likewise, with writing to a file there is a special built-in class in most programming languages that provides an interface for you. In the C-family this is called *StreamWriter* and in Python it is called *FileWriter*. Change the I/O type to *Write*. Take a look at the code extracts below.

Pseudo	
<pre> fileName ← directory address of file FileWriter ← new stream writer [fileName] For i ← 1 TO 5 Write "input line: " inputString ← Read written line FileWriter Writes inputString to file Next Close filewriter </pre>	<pre> Dim fileOut As StreamWriter Dim test As String Dim FileOut As StreamWriter fileName = "C:\MyDocuments\WritingToFileTest.txt" FileWriter = New StreamWriter(fileName) For i = 1 To 5 Console.WriteLine("input line: ") inputString = Console.ReadLine() FileOut.WriteLine(inputString) Next FileOut.Close() </pre>
C#	
<pre> Static StreamWriter FileWriter Static void Main (...) { String fileName = @"C:\MyDocuments\WritingToFileTest.txt"; fileWriter = new StreamWriter(fileName); for (int i = 1; i<= 5; i++) { Console.WriteLine("input Line Number {0}: ", i); String inputString = Console.ReadLine(); FileWriter.WriteLine(inputString); } FileWriter.Close(); } </pre>	<pre> fileName = "C:\MyDocuments\WritingToFileTest.txt"; file = Open(fileName, FileMode.Create); For x in 1 to 5 inputString = Console.ReadLine() file.WriteLine(inputString) file.close() </pre>
Pascal/Delphi	
<pre> var fileOut:textfile; inputString:String; count:integer; begin assignfile(fileOut,'C:\MyDocuments\WritingFileTest.txt'); reset(fileOut); // sets the program ready to write for count:=1 to 5 do begin readln(inputString); writeln(fileOut,inputString); end; closefile(fileOut); end. </pre>	

Binary files

Binary files are stored as binary encoded data. The content of the file itself is irrelevant but can be written in binary (see Section 5.2), so the content seems irrelevant. It can be read directly by components without the need for translating or conversions. Binary record types and read back into that type.

INSPECTION COPY

COPYRIGHT
PROTECTED



Writing binary files using record structure

We can declare a record, input the appropriate data and write it to the file in a format other than simple text which has to be converted.

Pseudo	
<pre> Declare recStructure as record name: string houseNum : integer postcode: string CurrentRec ← recStructure pathName ← #file pathway open binaryfile using pathname Set BinaryFile for Write Loop Input currentRec BinaryWriter (currentRec) Until finished or error in data Close binaryfile </pre>	<pre> Public Structure recStructure Public name As String Public houseNum As Integer Public postcode As String End Structure Dim CurrentRec As recStructure Dim CurrentFileReader As FileStream Dim CurrentFile As FileStream Dim Filename As String Filename = # file pathway CurrentFile = New FileStream(Filename, FileMode.Create) CurrentFileWriter = New StreamWriter(CurrentFile) Do CurrentRec.name = Console.ReadLine() CurrentRec.houseNum = Console.ReadLine() CurrentRec.postcode = Console.ReadLine() CurrentFileWriter.WriteLine(CurrentRec.name) CurrentFileWriter.WriteLine(CurrentRec.houseNum) CurrentFileWriter.WriteLine(CurrentRec.postcode) Console.WriteLine("do you want to add another record?") answer = Console.ReadLine() Loop until (answer = "N") CurrentFileWriter.Close() CurrentFile.Close() </pre>
C#	Pascal
<pre> Struct recStructure { Public string name Public int housenum Public string postcode } Static currentRec recStructure; Static BinaryWriter currentFileWriter; Static FileStream currentFile; String filename = # file pathway; currentFile = new FileStream(filename, FileMode.Create); currentFileWriter = new BinaryWriter(currentFile); do { CurrentRec.houseNum = Console.ReadLine(); CurrentRec.postcode = Console.ReadLine(); CurrentRec.name = Console.ReadLine(); currentFileWriter.Write(CurrentRec.houseNum); currentFileWriter.Write(CurrentRec.postcode); Console.WriteLine("do you want to add another record?"); answer = Console.ReadLine(); } While (answer == "Y"); currentFileWriter.Close(); currentFile.Close(); </pre>	<pre> type recStructure = record name:string[15]; housenum:integer; postcode:string[15]; end; var count:integer; fileOut:file of recStructure; currentRec:recStructure; answer:string; begin assignfile(fileOut, #file pathway, fmCreate); rewrite(fileOut); repeat readln(currentRec.name); readln(currentRec.housenum); readln(currentRec.postcode); write(fileOut, currentRec); writeln('Do you want to add another record?'); readln(answer); until answer = 'N'; closefile(fileOut); end. // note that in Pascal, the record structure is declared in the record type </pre>

INSPECTION COPY

COPYRIGHT
PROTECTED



Reading binary files using record types

By using record types and binary files the records are written to the file in the structure. However, when reading the files it is essential the same structure is used.

The examples below show a simple method with no error detection.

Pseudo	
<pre> Declare recStructure as record name: string houseNum : integer postcode: string CurrentRec ← recStructure pathName ← #file pathway open binary file using pathName binaryReader ← BinaryReader (recStructure) current ← binaryReader OUTPUT current </pre>	<pre> Public Structure recStr Public name As String Public houseNum As Integer Public postcode As String End Structure Dim CurrentRec As recStr Dim CurrentFileReader As BinaryReader Dim CurrentFile As FileStream Dim Filename as String Filename = # file pathway CurrentFile = New FileStream(Filename, FileMode.Open) CurrentFileReader = New BinaryReader(CurrentFile) Do While CurrentFile.Position < CurrentFile.Length CurrentRec.name = CurrentFileReader.ReadString() CurrentRec.houseNum = CurrentFileReader.ReadInt32() CurrentRec.postcode = CurrentFileReader.ReadString() Console.WriteLine(CurrentRec.name) Console.WriteLine(CurrentRec.houseNum) Console.WriteLine(CurrentRec.postcode) Loop CurrentFileReader.Close() CurrentFile.Close() </pre>
C#	
<pre> Struct recStructure { Public string name Public int housenum Public string postcode } Static currentRec recStructure; Static BinaryReader currentFileReader; Static FileStream currentFile; String filename = # file pathway; currentFile = new FileStream(filename, FileMode.Open); currentFileReader = new BinaryReader(currentFile); do { CurrentRec.name = currentFileReader.ReadString(); CurrentRec.houseNum = currentFileReader.ReadInt32(); CurrentRec.postcode = currentFileReader.ReadString(); Console.WriteLine(CurrentRec.name); Console.WriteLine(CurrentRec.houseNum); Console.WriteLine(CurrentRec.postcode); } While (currentFile.Position < currentFile.Length); currentFileReader.Close(); currentFile.Close(); </pre>	<pre> type recStr = record name: string; housenum: integer; postcode: string; end; var count: integer; filename: string; currentFile: FileStream; currentFileReader: BinaryReader; begin assign(filename, # file pathway); reset(currentFile, FileMode.Open); while not currentFile.Position = currentFile.Length do begin read(currentFileReader, count); write(currentRec, currentFileReader); write(currentFile, currentRec); write(currentFile, count); end; close(currentFile); end. </pre>

Note: # file pathway is the full path filename, for example 'C:\mydocuments\file.txt'

INSPECTION COPY

COPYRIGHT
PROTECTED



Another method of reading binary files is to use streaming to read/write data. Binary files whose contents can be used directly by hardware processors or the central processor if the processor knows how the data is formatted.

This processing of binary files is called streaming as the data is 'streamed' or 'streamed' over a program. The method is shown below.

Reading binary files with no record structure (streaming)

Here is the code that can be used for reading binary files.

Pseudo	
<pre> Current ← Null pathName ← #file pathway readStream ← FileStream (pathName, OpenFileMode) binaryReader ← BinaryReader (readStream) current ← binaryReader.ReadString() OUTPUT current </pre>	<pre> FileStream readStream; BinaryReader binaryReader; String current = null; String pathName = #file pathway; readStream = new FileStream (pathName, FileMode.Open); binaryReader = new BinaryReader (readStream); current = binaryReader.ReadString(); Console.WriteLine(current); </pre>
C#	
<pre> FileStream readStream; String current = null; String pathName = # file path way readStream = new FileStream (pathName, FileMode.Open); binaryreader = new BinaryReader (readStream); current =binaryReader.ReadString(); Console.WriteLine(current); </pre>	<pre> file = open(pathName, FileMode.Open); try: byte = file.ReadByte(); while byte != -1: # Convert byte to string # ... byte = file.ReadByte(); finally: file.Close(); </pre>
Pascal/Delphi	
<pre> uses classes,sysutils; var fsOut : TFileStream; fsIn : TFileStream; source: array[0..4] of integer = (2, 1, 8, 6, 244); begin fsIn := TFileStream.Create('binaryfile.bin', fmOpenRead); fsIn.Read(source, sizeof(source)); fsIn.Free; // the array Source now contains the data end. </pre>	

INSPECTION COPY

COPYRIGHT
PROTECTED



Writing binary files with no record structure (streaming)

Writing binary files is similar to writing text files and you will notice the similarity in the family languages. In this example you will see that an array containing some values is being written to the file.

Pseudo	
<pre>Source ← #values for array Using writer ← #open file, filename, file mode For i ← 1 to source.length Print</pre>	<pre>Dim source[5] Using writer As BinaryWriter = File.Open(filename, FileMode.Create) For value = 0 To source.Length - 1 Writer.Write(source[i]) Next End Using</pre>
<pre>int source[5] = {2, 1, 8, 6, 244}; Using BinaryWriter = New BinaryWriter (file.Open("binaryFile.bin", FileMode.Create)) { For (int value = 0; i <= source.Length; i++) { writer.write(value); } }</pre>	<pre>with open('binaryFile.bin', 'wb') as f: for value in source: f.write(value)</pre>
Pascal/Delphi	
<pre>uses classes,sysutils; var fsOut : TFileStream; source: array[0..4] of integer = (2, 1, 8, 6, 244); begin fsOut := TFileStream.Create('binaryfile.bin', fmCreate); fsOut.Write(source, sizeof(source)); fsOut.Free; // this prevents memory leaks end.</pre>	<p>Task:</p> <p>Write the code to write the array to the file. The code should be able to write the array to the file. When the file is closed, all the data should be displayed.</p>

INSPECTION COPY

COPYRIGHT
PROTECTED



ABSTRACT DATA TYPES / DATA STRUCTURES

Abstract data types are defined as data types which are not defined by their programming language. Instead they are defined by the operations that can be performed on them.

The abstract types and structures that you need to be familiar with are as follows:

- Queues
- Stacks
- Lists
- Graphs
- Trees
- Hash tables
- Dictionaries
- Vectors

Each of these is covered in detail over the following pages.

2.2 QUEUES

The data structure known as a queue has the same characteristics as the queue you encounter in everyday life. For instance, a queue at the checkout counter in a supermarket increases at its rear as customers join the queue to have their purchases checked out, and only reduces in size when a customer is served at the front of the queue, the checkout counter. A queue of cars at traffic lights behaves in a similar manner, with cars exiting the queue only at its front and joining the queue only at its rear. This is a FIFO data structure (First In, First Out).

A queue requires two pointers, one of which points to the front and the other to the rear. These pointers need to be set up in such a way that the following operations can be carried out:

1. Check whether the queue is empty.
2. Return the value of the first element (front).
3. Return the value of the first element and remove it.
4. Place new element onto the rear of the queue.


There are three different types of queue which you need to know about: circular, priority and dequeues. Circular queues are particularly suited to implementation as arrays, and lend themselves to being implemented with lists.

CIRCULAR QUEUES


A circular queue is a queue which has a fixed amount of space, but where the front and rear pointers can wrap around, much like a circle. This structure lends itself easily to buffering data.

1. The *front* points to the element of the array which should be removed.
2. The *rear* points to the last element added. As data is added these pointers move along the array and loop back to the start of the array.

Suppose a queue is formed in the following order: Ashley, George, Julie, Jacques. An array as follows (as an array of 1 by 10 array):




Suppose that two names (Ashley and Julie) leave the queue and two new names (Justin and Mary) join the queue. The queue would now look like this:



COPYRIGHT
PROTECTED



As more names are added the array would eventually loop round and if the Kumar were added (in that order), the queue would look like this:



If another name were to be added to this queue then the queue would be full if the front and tail are next to each other, with the front on the right of the array to the first element in the array and the tail points to the last element in the array. To check that the queue is empty by looking at the tail and checking if it is equal to the front.

Implementation of a circular queue

We are assuming the array is of size 10 (i.e. data [10]). This has the limitation of a maximum of 10 items. We need two pointers that point to the front and the rear.


test for an empty queue:	To look at the front of the queue:
<pre> PROCEDURE see_empty() IF front < rear THEN PRINT("Queue is empty!") ELSE PRINT("Queue is not empty") END IF END PROC </pre>	<pre> PROCEDURE see_front() IF front < rear THEN PRINT("Queue is empty") ELSE PRINT("Front is " & data(front)) END IF END PROC </pre>
To add (push) an item to a queue: (assumes the pointers and array are global variables)	To take an item from the queue: (The function will return the item)
<pre> PROCEDURE push(new_item) IF (rear + 1 < front) OR (rear < 10 AND front < 1) THEN PRINT("Queue is full!") ELSE IF rear < 10 THEN rear < 10 ELSE rear < rear + 1 END IF data(rear) < new_item END IF END PROC </pre>	<pre> FUNCTION pop() IF front < rear THEN PRINT("Queue is empty") ELSE pop < data(front) IF front = 10 THEN front < 1 ELSE front < front + 1 END IF END IF END FUNCTION </pre>
To print out all the items in a queue	
<pre> PROCEDURE print_queue() DIM i As Integer // loop counter DIM queue_string As STRING //builds up the queue before printing i < front WHILE i < rear queue_string < queue_string & data(i) IF i < 10 THEN i < 1 ELSE i < 10 i < i + 1 End WHILE PRINT(queue_string) END PROC </pre>	

**COPYRIGHT
PROTECTED**




LINEAR QUEUES


Since a queue usually holds a bunch of items with the same type, it makes sense to store it in an array. With linear queues, elements are always added at one end and removed from the other. The diagram below demonstrates the concept using a fixed-size queue of 8 elements (front and rear pointers are at the ends of the queue).



Items are always removed from the front of the queue. To do this, the front pointer is moved to the next item in the list. The queue below shows the element (Apple) being removed.



Items are added to the list at the rear end. The example below shows two new items added to the list. As with the front pointer while adding elements, the rear pointer is moved so that it points at the newest item:



PRIORITY QUEUES

Priority queues are queues where items are removed in order of their priority. The items are added to the queue in any order, but they are removed in order of their priority. Each item is assigned a priority as it is added to the queue. The priority can be based on the basis of origin, data type, time of day – anything really.

Priority queues find uses in a number of areas of computing. Network buffers use priority queues. By using priority queues in networks, it is possible to move important communications, to the front of the queue, while at the same time moving less important communications, to the back of the queue. For example, file-sharing applications, to the back of the queue.

Adding an item	Removing an item
<pre> IF queue is full Then OUTPUT Error ELSE rearPointer++ IF rearPointer > maxIndex Then rearPointer = 1 END IF queue[rearPointer] ← datum END IF </pre>	<pre> IF queue is empty Then OUTPUT error ELSE Return queue[frontPointer] frontPointer++ IF frontPointer > maxIndex Then frontPointer = 1 END IF END IF </pre>
Testing if empty	Testing if full
<pre> IF queue [frontPointer] ← 1 Then Return true ELSE Return false END IF </pre>	<pre> IF queue[rearPointer] = maxIndex Then Return True ELSE Return False END IF </pre>

COPYRIGHT
PROTECTED



Questions: Queues

- 1 You are tasked with writing a video buffer for a video player that can store a limited number of frames. Which would be more suitable, a stack or a queue? Explain your answer. (2 marks)
- 2 A circular queue is usually implemented as an array.
 - a) What variables are required to keep track of such a queue? (1 mark)
 - b) How can you check whether a circular queue is empty without using a counter? (1 mark)
 - c) How can you check whether a circular queue is full without using a counter? (1 mark)
 - d) Write a procedure to add an item to the queue in pseudocode. Include an error message if the item cannot be added. (1 mark)
- 3 Consider the following linear queue. Elements are added at the rear and removed from the front.

FRONT	REAR				
0	1	2	3	4	5
SF	A2	EE	72		

- a) Draw the queue after two elements have been removed. (1 mark)
- b) Devise a method for representing an empty queue. What steps would you need to take for push/pop procedures? (1 mark)
- c) Write a procedure, using pseudocode, to add items to the queue when the queue is empty. (1 mark)

2.3 STACKS


A stack is a data structure characterised by the expression Last In, First Out (LIFO). The most recent item added to the stack is the first one which can be removed from the stack. You must keep track of the last item added to the stack – that is, the current top of the stack.

A real-life visualisation of a stack is the stack of trays at the entrance to a canteen. If you need to take that tray off the top in order to get to the next one down. If you have finished with that tray and put the dirty tray back on the top of the stack. However, you cannot take a tray from the bottom of the stack without first removing all the other trays!

A stack needs to be set up so that the following operations can be carried out:

- Check whether the stack is empty (NULL or not)
- Check whether the stack is full (not)
- Look at the top value and remove it (pop)
- Look at the top value, without removing it (peek)
- Insert a new value on top of the existing stack (push)

Here is an example of a stack used to hold names, and how pushing items on and removing items from the stack works:



COPYRIGHT
PROTECTED



Implementing a stack

A stack can be implemented using an *array* or a *linked list*. To implement a stack using an array, a variable is needed to keep track of the top of the stack.

For example, an array declared like the following in Java could be used as a stack. A variable is needed to keep track of the top of the stack and decrementing it when an element is removed.

```
String[] myStack;
myStack ← new String[10];
int top ← -1;           // -1 means the stack is empty
```

Alternatively a linked list can be used instead of an array, with only the address of the next element needed.

Example – Procedures to implement a stack using arrays

What follows are two examples of procedures, written in pseudocode, which implement a stack using an array.

We are making the following assumptions:

1. There is an integer variable, *top*, which points to the top of the stack. If *top* is -1, the stack is empty.
2. The size of the array *myStack* is given by the integer *max*.
3. Elements are accessed in the array by reference, with 0 being the first element.
 - *myStack[0]* is the first element
 - *myStack[1]* is the second element
 - *myStack[299]* is the three-hundredth element
 - *myStack[max-1]* is the last element

Testing an empty stack is simple as all we need to do is check whether the top of the stack is -1.

```
PROCEDURE test_empty( )
  IF top = -1 THEN
    PRINT("Stack is empty")
  ELSE
    PRINT("Stack is not empty")
  END IF
END PROCEDURE
```

To look at the top item on a stack without removing it, we can say:

```
PROCEDURE print_top( )
  IF top = -1 then
    PRINT("Stack is empty")
  ELSE
    PRINT("Top item on stack is " & myStack[top])
  END IF
END PROCEDURE
```

**COPYRIGHT
PROTECTED**



To add an item (push) to a stack we can use the following procedure (assuming variables):

```
PROCEDURE push(newItem)
  IF top = max - 1 THEN
    PRINT("Stack is full!")
  ELSE
    top ← top + 1
    myStack[top] ← newItem
  END IF
END PROCEDURE
```

To take off an item (pop) from a stack we can use the following function (which returns the item):

```
FUNCTION pop( ) AS TYPE OF myStack
  IF top = -1 THEN
    PRINT("Stack is empty!")
  ELSE
    top ← top - 1
    RETURN myStack[top]
  END IF
END FUNCTION
```

Questions: Stacks

- 1 What rule is said to govern stack data? (1 mark)
- 2 Imagine a stack containing the following numbers: 89, 45, 22, 90
90 is the top of the stack.
 - a) Rewrite the stack after the following operations have been performed:

```
pop()
pop()
push(77)
push(56)
```

- b) Write a function in pseudocode that will add up all the items in the stack.
- 3 Here is an example of a stack implemented as an array:

Top of stack: 5

0	John
1	Lara
2	Mike
3	Steve
4	Frank
5	Billy
6	Georgina


- a) Which name would be removed from the stack first when the pop operation is performed?
- b) What steps are needed to push an item onto the stack? (1 mark)
- c) What would be a suitable value for the top variable when the stack is empty?

**COPYRIGHT
PROTECTED**




2.4 GRAPHS

A graph has a set of *vertices* (often also referred to as *nodes*) and *edges* where vertex is a point and edges are the lines that join the points together. A labelled graph has vertices associated with a label.




Directed graph

A directed graph (also known as digraph) consists of arcs and vertices where an arc is a directed edge and has a direction. A labelled digraph is one that has its vertices associated



Undirected graph

Undirected graphs can be represented using directed graphs very easily by simply replacing each undirected edge with two edges, each pointing in opposite directions.




COPYRIGHT
PROTECTED



Weighted graphs

Weighted graphs are graphs where every edge is given a weight. The weight represents some quantity. A very common use of weighted graphs is to represent locations. The travelling salesman problem is an example of a problem where the weights represent distances. However, the weights could represent anything; for example, they could represent the cost of a link when analysing network performance or they could represent the number of links between two web pages in a search engine.

Below is an example of a weighted undirected graph representing the distances between four cities in the travelling salesman problem:




The uses of graphs

Graphs can be used in many applications, for purposes such as finding routes, designing computer networks, traffic control and finding the best route to take in solving games such as mazes and Sokoban puzzles. The trick is being able to represent the problem by using a graph.

Adjacency list

An adjacency list is used to represent either edges or arcs. It is a linked list where each node contains the current node, and what nodes the current node is connected to. Consider the following graph:



The adjacency list would be as follows:

Vertex	Connected to
A	B, D
B	C
C	A
D	C

**COPYRIGHT
PROTECTED**




Adjacency matrix

An adjacency matrix is another method to represent edges in a graph. An n -by- n is used to represent all the edges in the graph where n is the number of vertices

The adjacency matrix shows how many edges are used to connect each vertex together. For example, vertex A is connected with vertex B by one edge.

To come back to the travelling salesman problem a weighted adjacency matrix can be used.




	Bristol	Liverpool	London	Plymouth
Bristol	0	181	119	118
Liverpool	181	0	0	0
London	119	0	0	238
Plymouth	118	0	238	0
Sheffield	179	78	166	0

In the above the weighting between Bristol to Liverpool is 181 however if it was from Bristol \rightarrow Liverpool, and not from Liverpool \rightarrow Bristol, then the Liverpool

Questions: Stacks

- 1 Consider the following directed graph:




- Write an adjacency list to represent this digraph. (1 mark)
- Write an adjacency matrix to represent this digraph. (1 mark)
- Assuming space is the main limiting factor, which would be the why? (2 marks)

COPYRIGHT
PROTECTED




2.5 TREES

A tree is a simple undirected graph that contains no cycles in it, i.e. each vertex is part of only one path. A tree with n vertices always has $n - 1$ edges.



ROOTED TREE

A rooted tree is usually used to show a type of hierarchy. It contains a root vertex and a stem.




Generally rooted trees are drawn with the root at either the top or bottom of the diagram easier to understand and makes the structure clearer. Any vertex can be designated as such.

BINARY TREES

A binary tree is a type of rooted tree which is often used in computing as an efficient way of searching. Because they are more structured and therefore quicker to search than arrays, they can be used as a basis for constructing a database.

Binary trees consist of *nodes* (vertices) and *branches* (edges), which are the links between nodes. Nodes which have no children are called *leaf* nodes. A binary tree must always be used for inserting and deleting elements or else it would be pointless to use one since searching for an element in a node consists of a field and two pointer values, one for the left subtree and the other for the right subtree.



Binary trees are a form of abstract data structure and therefore, like linked lists, can be implemented in a number of ways. Binary trees can be implemented as arrays with one column for the data and two pointers (left and right). They can also be implemented dynamically using pointers.

INSPECTION COPY

COPYRIGHT
PROTECTED



Constructing a binary tree

By convention, items added to a binary tree are compared to the root node; if greater than the root node they are placed on the right-hand side and if less than the root node they are placed on the left-hand side. This is then repeated for every node they come across until they reach a position where they can be added.

A procedure to produce a binary tree would therefore need to contain the following steps:

1. Create a root node.
2. Insert data into the root node.
3. If the root node exists then compare the new item to be added with the current item as the root node.
4. If item is > root node follow right pointer, if not follow left pointer.
5. Compare until the left or right pointer is Null. That is the appropriate position to add the item.
6. Create node, insert item into node and set both left and right pointers to Null.
7. Connect the new node to the tree.

Adding an item to a tree

To add an item to a tree we can use the following procedure (assuming *array* variables and the first element is always in row 1):

```

PROCEDURE add_leaf(new_item)
    row ← 1
    current ← 1
    finished ← false

    WHILE (row <= arraysize and tree(row, 1) != "") // loop until we find a null pointer
        row ← row + 1
    END WHILE

    IF row > arraysize // check that array is not full
        PRINT("Tree is full!")
    ELSE
        tree(row,1) ← new_item // puts new item in
        tree(row,2) ← -1 // sets pointers to null
        tree(row,3) ← -1 // sets pointers to null
    END IF

    WHILE finished = false
        IF new_item < tree(current,1,1) // adds pointer to new item
            IF tree(current,1,2) = -1
                tree(current,1,2) ← row
                finished ← true
            ELSE
                current ← tree(current,1,2)
            END IF
        ELSE IF tree(current,1,3) = -1 THEN // adds pointer to new item
            tree(current,1,3) ← row
            finished ← true
        ELSE
            current ← tree(current,1,3)
        END IF
    END WHILE
END PROCEDURE


```

**COPYRIGHT
PROTECTED**




The following diagram shows how the first six steps are used to construct the

Step 1 – Add Flavia



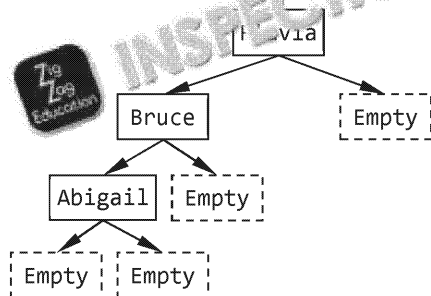
Flavia is the first node added and so becomes the root of the tree.

Step 2 – Add Bruce




Bruce is before Flavia so go left.
Empty position found so Bruce is added here.

Step 3 – Add Abigail




Abigail is before Flavia so go left.
Abigail is before Bruce so go left.
Empty position found so Abigail is added here.

Step 4 – Add Millie




Millie is greater than Flavia so go right.
Empty position found so Millie is added here.

Step 5 – Add Serena



Serena is greater than Flavia so go right.
Serena is greater than Millie so go right.
Empty position found so Serena is added here.

Step 6 – Add Rachel



Rachel is greater than Flavia so go right.
Rachel is greater than Millie so go right.
Rachel is less than Serena so go left.
Empty position found so Rachel is added here.

Questions: Trees

1. Consider the set of integers $a = \{7, 2, 6, 11, 5, 9, 4, 8\}$. If the items, before being inserted into a binary tree, were sorted into descending order what would the tree look like?
2. The search time for a binary tree is usually $O(\log n)$. If the binary tree search time became $O(n)$, what would the search time become? (1 mark)
3. Describe a procedure to delete a node in a binary tree implemented using an array.

**COPYRIGHT
PROTECTED**



2.6 HASH TABLES

So far you have looked at algorithms which can be used to search and sort arrays. One problem with all of these structures is that as data volume increases in arrays becomes more difficult and harder to manage. All of these structures require a structure like a tree, which allows a binary search operation, increases in time. For more detail, refer to the notes on Big O Notation (Section 4.4). A binary search

The hashing algorithm is a mathematical calculation performed on search criteria. If we were to use the surname we could perform a calculation to find the data location (where the data to be found). If well devised it can significantly cut down the search time in the worst case; often it is $O(1)$, i.e. one iteration).

Hashing algorithms are often tailored to the specific application that they will be used for. One hashing algorithm which is used a lot is the FNV-1a. One example of a hashing algorithm operates on the first two letters of a name, where the letters a–z are given the values 1–26. The index is then produced by multiplying the first number by 26 and then adding the second number. Here is an example of a hash table showing keys that the algorithm will produce.

Surname	Jones	Zheng	Patel	Ba
Calculation	$9 \times 26 + 14$	$25 \times 26 + 7$	$15 \times 26 + 0$	$1 \times 26 + 0$
Index	248	657	390	26

As can be seen, if we wanted the information about the person whose surname was Patel, we would look at record location 390. So we could jump directly to memory or index 390 and find the information.

Like all hashing algorithms, however, there is a problem, and this is that it produces the same index for different surnames that begin with the same two letters. For example, the surnames Patel and Ba would have the same index. This is known as a collision. The aim of a good hashing algorithm is to avoid collisions while also optimising space needed.

Collisions

One of the main problems that hashing functions have is that they will probably produce the same index, i.e. two different items could produce the same index. This is called a collision. One collision resolution strategy is to use the next available index. This seems a good idea in principle but leads to problems later. For example, if a record is not in the position you expect, you can't assume it has been deleted because it might have been moved else due to a collision. Likewise, if you delete a record you have to indicate that you know there may have been a collision.

The solution to this is to use overflow lists. There are two possibilities for an overflow list.

Firstly, that all records with duplicate index values are simply put into another list. If a record cannot be found in the main table, the program will need to perform a linear search on that list. Only once the overflow list has been searched can the program ensure that the record is not in the table.

The other possibility is to have linked lists attached to every index. All the records with the same index are stored in the linked list, and the program simply performs a linear search on the linked list for the index it has calculated. However, this solution is slightly more complex to implement.

Questions: Hashing

- 1 Why is hashing important?
- 2 Hashing is often used to store data in binary form. Explain how it can be used to store data in binary form.
 - a) Work out the index for the surname *Hull, Brighton and Manchester*.
 - b) Place the place names in a list based on the index.
 - c) What problem does this cause if you add *Bristol* to the list?
 - d) Propose a solution to this problem without changing the index.

**COPYRIGHT
PROTECTED**



2.7 DICTIONARIES

Dictionaries are another example of an abstract data type which a programmer can use to store data in an ordered manner, although it is worth noting that it is also defined as a *generic class* in C programming.

If you take the analogy of a real-world dictionary, the data is stored in a two-part key: the *word* (the key) and the *definition* (the data). To get to the data you search for a specific word and read the definition; likewise in programming you search the key-field and retrieve the data associated with the key.

The functions that can be applied to a dictionary construct are dependent on are using, but all have the same basic functionality. For example, in the C *for* *generic class* and implemented using a list, therefore the class has all the functions included in functions of the dictionary that have been built into the language.

Creating and adding to a dictionary

Pseudo	
<pre>STRUCTURE dictionary (key, data) Dictionary[2] ← "Computer" Dictionary[1] ← "I" Dictionary[0] ← "Love" Dictionary[-1] ← "Science" END STRUCTURE</pre>	<pre>Dim dictionary As Integer) Dictionary Dictionary Dictionary Dictionary</pre>
C#	
<pre>Dictionary<string, int> dictionary = new Dictionary<string, int>(); dictionary.Add("Computer", 2); dictionary.Add("I", 1); dictionary.Add("Love", 0); dictionary.Add("Science", -1);</pre>	<pre>Dictionary = { } Dictionary["Computer"] Dictionary["I"] = Dictionary["Love"] Dictionary["Science"]</pre>
Pascal/Delphi	
<p>No native dictionary system is available in Pascal/Delphi, although there is a method utilising a library called Generics; the example below is from free Pascal using the library fgl:</p> <pre>uses fgl, sysutils; // using the generics library type TDictionary = specialize TFMMap<string, integer>; var i: integer; dictionary: TDictionary; begin dictionary := TDictionary.Create; dictionary.Sorted := True; // for fast lookup of keys // assign values dictionary['Computer'] := 2; dictionary['I'] := 1; dictionary['Love'] := 0; dictionary['Science'] := -1;</pre>	

INSPECTION COPY

COPYRIGHT
PROTECTED



Using a dictionary

Dictionaries are useful whenever you need to store or retrieve unique information. In an initialised dictionary there are a few functions you're going to want to know.

Checking whether an entry exists

If you want to check whether there is an entry in the dictionary using a key, you can output the value by using the following code:

Pseudo	
<pre>If dictionary Contains ["science"] Then Value ← dictionary "science" PRINT value ELSE PRINT "Value not found" END IF</pre>	<pre>If dictionary.ContainsKey("science") Dim int value = dictionary["science"] Console.WriteLine(value) Else Console.WriteLine("value not found") End If</pre>
C#	
<pre>If (dictionary.ContainsKey("science")) { Int value = dictionary["science"]; Console.WriteLine(value); } Else { Console.WriteLine("value not found"); }</pre>	<pre># method one: Print (dictionary["science"]) # method two: If "science" in dictionary: Print (dictionary["science"]) Else: Print ("value not found")</pre>
Pascal/Delphi	
<pre>if dictionary.find('Science',i) then writeln(dictionary['Science']) else writeln('value not found');</pre>	

Note: when using a member of the C family, you can also use the extension 'ContainsKey'. A better method you could research is method called 'TryGetValue'.

Removing an entry from a dictionary

Just as easily as adding an item to a dictionary, you can remove items too. This is done by removing redundant information, such as when you save over a file.

C# and VB.NET	Python	Python
Dictionary.Remove("Science");	del dictionary["Science"]	del dictionary["Science"]

Implementation without libraries

A dictionary may also be implemented using an array and utilising search routines. One method would be a linked list array-type structure. Another structure is illustrated in pseudocode below.

Initialising the dictionary:

```
NEW STRUCTURE ← TDictionary # declare the structure
    Item ← string
    Index ← integer
END STRUCTURE

Dictionary ← array of TDictionary
```

COPYRIGHT
PROTECTED



Using the dictionary (finding)

This method utilises a method of knowing how many items are in the array (

```

pointer ← 0
While pointer ≤ numOfItems or found = false
  If dictionary[pointer].Item = itemToBeFound then found = true
  Pointer ← pointer + 1
End while

if found then
  Output dictionary[pointer-1].index
Else
  Output "item not found"
Return pointer

```

Removing an item

This is a shuffling-up method as we are using a numOfItems. Other solutions would move the pointers.

```

Find(itemToBeFound)
IF itemFound then
  Shuffle array up from pointer+1 to numOfItems
  numOfItems ← numOfItems - 1

```

**COPYRIGHT
PROTECTED**



2.8 VECTORS


In mathematical terms, a vector is a geometric quantity for a position in space with direction but no location. Vectors can be expressed in a number of ways:

- As a list of numbers written in square brackets, e.g.:
 $[3.14159, -0.3, 8.0211127, 1.0]$
- As a power of a set where the power is the number of entries from the set
 \mathbb{R}^4
- As a function using a dictionary to map values to a set (' \mapsto ' means 'maps to')

$$\begin{aligned} 0 &\mapsto 3.14159 \\ 1 &\mapsto -0.3 \\ 2 &\mapsto 8.0211127 \\ 3 &\mapsto 1.0 \end{aligned}$$


A vector can also be depicted using 'arrow' notation. This is the easiest way to visualise the vector as a movement; the length of the arrow is the magnitude. You can think of a vector by its horizontal (x) and vertical (y) displacements.

For example, the vector below would be $[a, b]$:



If you're finding vectors hard you might be trying to imagine them in terms of a movement. Try to remember that although vectors have magnitude and direction, they don't have a location. What makes them so useful.

Take a look at the example below.



Both
have
This
have

**COPYRIGHT
PROTECTED**




Vector addition and subtraction

Look at the following example; note how the two vectors when used together the hypotenuse is given by the form $[a + b, c + d]$ – this transformation is called using a vector to contain the magnitude and direction of the movement.

The translation can be given by the equation: $[a, b] + [c, d] = [a + c, b + d]$ where


The equation becomes $[a, c] + [b, d]$ because you group similar transformation transformations in the same direction and their magnitudes are added together.



It is easier to visualise this with an example of moving a shape in two-dimensions. Translating the triangle given by the vector $[a, b]$ by $[j, k]$.

Before translation

After translation



The values of these points are arbitrary but you can clearly see how the translation is performed. The distance and the direction of the transformation are given by the values $[a, b]$ and $[j, k]$. The equation for the translation is: $[a, b] + [j, k] \rightarrow [a+j, b+k]$.


COPYRIGHT
PROTECTED



Vector multiplication

This transformation is nice and easy. You will see later in *Vector Graphics* (see page 28) that it is perfect for representing objects where high precision is needed because they can be moved, inverted and scaled and they always look the same. Vector multiplication comes in – it achieves *scaling*.

Consider the example in the introduction where you represented a vector with its components and labelled them 'a' and 'b'. This will highlight how simple it is to scale a vector using its scalar product. If you multiply each value of the vector by a scalar value, you will result in a vector of twice the scale. $2[a, b] \rightarrow [2a, 2b]$



As you can see, the direction of the vector is maintained and the magnitude is doubled.

Dot products of vectors (vector scalar product)

Consider multiplying two vectors together that have the form:

$$u = [u_1, u_2, \dots, u_n], v = [v_1, v_2, \dots, v_n]$$

The result can be written as:

$$u \cdot v = u_1v_1 + u_2v_2 + \dots + u_nv_n$$

This can be seen when cross-multiplying matrices, such as in the following example:

$$\begin{bmatrix} k & l \\ m & n \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} ka + lc & kb + ld \\ ma + nc & mb + nd \end{bmatrix}$$

You can see the scalar product has been used: $ka + lc = [k, l] \cdot [a, c]$

Questions: Vectors

1 $B = \begin{bmatrix} K \\ 0 \end{bmatrix}$



Based on the above, what is the name of the transformation represented by B?

2 $A = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$

Based on the above, what is the vector notation of the transformation represented by A?

3 $A = \begin{bmatrix} 3 \\ 1 \end{bmatrix}$ and $C = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$

Based on the above, consider the transformation required for the addition of A and C. What is the vector notation for this transformation? (2 marks)

COPYRIGHT
PROTECTED



3. Algorithms

This section covers the fundamentals of algorithms, illustrated by common examples. Algorithms are often used in examinations, therefore it is important to be able to reproduce, trace and have a strong grasp of the concepts.

This section covers:

3.1 Graph traversal.....	p1	3.4 Searching algorithms.....	p1
3.2 Tree traversal.....	p6	3.5 Sorting algorithms.....	p6
3.3 Reverse Polish notation (RPN).....	p8	3.6 Dijkstra's shortest path.....	p8

3.1 GRAPH TRAVERSAL


Graph traversal algorithms are algorithms which visit all of the nodes in a graph. This is necessary for many operations, for example, searching for a particular node or more complex operations, such as graph colouring, which are not covered here.

Breadth-first and *depth-first* are two possible ways of traversing a graph.

- *Breadth-first* means going through the graph starting from a particular node, then looking at the nodes connected to those nodes.
- *Depth-first* means going through the graph, again starting from a particular node. The nodes connected to the first node connected to the start node are completely explored, then the second node connected to the start node and completely exploring all its children and so on.

The diagram below shows the order in which a graph would be traversed using breadth-first search (assuming nodes on the left are visited before nodes on the right).

Breadth-first Traversal



BREADTH-FIRST SEARCH

A breadth-first search first searches all the nodes which are adjacent to the start node, then the nodes connected to them. The informal algorithm is as follows:


1. Choose a vertex to start from.
2. Colour the start vertex (cross-hatch) to indicate that it has been visited.
3. Create a list of neighbouring nodes and visit one of them.
4. Colour the visited node (cross-hatch), then visit the other nodes in the list.
5. Repeat this process until no more nodes are left in the list.
6. Colour the start vertex black to indicate that it has been completely visited.
7. Go to the next vertex and repeat the process as above.

**COPYRIGHT
PROTECTED**



Performing a breadth-first search on a graph

Here is an example of the steps taken during a breadth-first traversal of the




INSPECTION COPY




indicates a node which has been **visited**

indicates the node has been **completely explored**


①




②



③






④




COPYRIGHT
PROTECTED



**COPYRIGHT
PROTECTED**



The order that the nodes have been traversed in is:



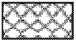

DEPTH-FIRST SEARCH


Depth-first search is another simple graph traversal search, which instead of depth-wise. The algorithm is as follows:

1. Select a start vertex and colour it (cross-hatch) to indicate that it has been visited.
2. Visit a new vertex adjacent to the current vertex that has not been visited.
3. Colour the visited vertex (cross-hatch) and visit a new vertex that is adjacent to it.
4. Repeat these steps until there are no more adjacent vertices to visit.
5. Colour the vertex black to show that it has been completely discovered.
6. Visit an adjacent vertex that hasn't been visited and colour it grey.
7. Repeat these steps until all the vertices are coloured black.
8. If there are any remaining vertices that have not been coloured due to being unreachable, set the start vertex to be one of the unreachable vertices and repeat the process until all vertices are reached.

Perform a depth-first search on a graph


Here is an example of the steps taken during a depth-first traversal of the following graph on the right.

-  indicates a node which has been **visited**
 indicates the node has been **completely explored**.




COPYRIGHT
PROTECTED






The order that the nodes have been traversed in is:



Note: depending on which node you pick, you may get a different output! How would the result change if you picked the right node instead of the left in Step 4?

Questions: C

1 Consider the



What order w
starting from

- a) A breadth
- b) A depth-

**COPYRIGHT
PROTECTED**



3.2 TREE TRAVERSAL

Searching a binary tree ($O(\log n)$)

Since everything is organised in a hierarchical manner, searching for an item is complicated. It is the same fundamental process as inserting an element. The steps are as follows:


1. Start from root node.
2. If item is equal to the item in the current node then item is found, then return the current node.
3. If item $>$ the current node, follow the right pointer.
4. If item $<$ the current node, follow the left pointer.
5. Repeat the process for the next node until a null pointer is found.
6. If Null pointer found, then the item is not in the tree (but could be added).

Other tree traversal algorithms

An algorithm that traverses a tree will perform an action on all nodes in the tree. The action could be to print the value of the node, even if it is only to compare the node to another to find the minimum or maximum. The three types of traversal are *pre-order*, *in-order* or *post-order*. Note the use of recursion in the procedures. Recursive algorithms are much more elegant. It is extremely hard to do these operations with iterative algorithms.

Suppose the following letters have been added to the first tree where letters smaller than the root node are stored to the left, and letters larger than the root node are stored to the right. The algebraic expression $A + B * C - D$ has been stored in the second tree with the operators to the left and right of the root node.

The trees would now look like this:



Suppose that each tree is stored in three arrays called data, leftPointer and rightPointer. The arrays would look like this:

	Data	leftPointer	rightPointer
1	K	2	3
2	D	5	4
3	T	-1	-1
4	D	6	7
5	M	-1	-1
6	A	-1	-1
7	F	-1	-1

	Data	leftPointer
1	*	2
2	+	4
3	-	6
4	A	-1
5	B	-1
6	C	-1
7	D	-1

Note: -1 means that the specific element has no children to that side.

COPYRIGHT
PROTECTED



Pre-order traversal

Pre-order is when the algorithm looks at the root node, then makes its way down the left-hand side and then the right-hand side. The first tree above would be listed in pre-order as $*+AB-CD$.

An algorithm for a pre-order traversal follows, where p is initially the array

```
PROCEDURE preOrder (p)
  Print (data[p])
  IF leftPointer[p] > 0 THEN
    preOrder(leftPointer[p])
  END IF
  IF rightPointer[p] > 0 THEN
    preOrder(rightPointer[p])
  END IF
END PROCEDURE
```

In-order traversal

In-order is when the algorithm first goes down the left branch, then looks at the root node and then the right branch. The first tree above would be listed in in-order as A, D, F, K, M, N. The second tree above would be listed as $A+B^*C-D$.

An algorithm for an in-order traversal follows:

```
PROCEDURE inOrder(p)
  IF leftPointer[p] > 0 THEN
    inOrder(leftPointer[p])
  END IF
  PRINT(data[p])
  IF rightPointer[p] > 0 THEN
    inOrder(rightPointer[p])
  END IF
END PROC
```

Post-order traversal

Post-order is when the algorithm first goes down the left branch, then goes down the right branch and then looks at the root node. The first tree above would be listed in post-order as A, D, F, K, M, N. The second tree above would be listed as $AB+CD-*$.

An algorithm for a post-order traversal follows:

```
PROCEDURE postOrder(p)
  IF leftPointer[p] > 0 THEN
    postOrder(leftPointer[p])
  END IF
  IF rightPointer[p] > 0 THEN
    postOrder(rightPointer[p])
  END IF
  PRINT(data[p])
END PROCEDURE
```

As you can see, the algorithms are remarkably similar. It may help you to remember which is which by considering the root node value before the left and right nodes (pre-order), after the left but before the right node (in-order) or after the left and right nodes (post-order).

Questions: Tree Traversal

- Construct a binary tree from the following numbers in the order they are given: 0 -1 1 2 -2. Write out the order in which the tree was traversed in:
 - Pre-order
 - Post-order
 - In-order

**COPYRIGHT
PROTECTED**



3.3 REVERSE POLISH NOTATION (RPN)

An arithmetic expression is a form of numbers and operators which represent a calculation. When we write down an arithmetic expression; these include the infix and postfix (Reverse Polish) notation.

Infix notation is what you all use every day. This is when operators and symbols are placed between the operands. An operator is surrounded by two symbols, one on the left and the other on the right. Examples of infix expressions include $5 + 5$, $2 * 10 + 2$, $16 / 4 + 3$. In general use, it requires that brackets be used to determine the order of operations. For example, $5 + 5 / 2$ is different to $5 + 5 / 2$.

In the early days of computing, it was very complicated to program computers to solve a problem. This problem, Reverse Polish notation was invented. This places the operator after the operands. As you will see, this eliminates the requirement for brackets.

In Reverse Polish notation operators are evaluated from left to right. So the expression $2 * (3 + 7) + 1$ can be written without brackets in infix without increasing the number of operations. The expression can be written as $2 * (3 + 7) + 1$. Notice that writing it with brackets does not change the order of operations.

Here are some more examples of infix expressions converted to postfix expressions. Note that in every case there are a variety of different orders the expression could be written in:

Reverse Polish notation has another advantage on top of allowing expressions to be written unambiguously without brackets.

Evaluating RPN using a stack

The stack can be used to evaluate Reverse Polish notation by the algorithm given below. A stack is a LIFO data structure; this means that the first thing that is put into the stack is the last thing that is removed.

- Go through the input and read one character at a time until there is no more input.
- Whenever a symbol is read from input then push (put) symbol on stack.
- Whenever an operator is read from input then pop (get) the last two symbols from the stack, perform the operation on those two symbols read and store the answer back onto the stack.

Example: $1\ 3\ 3\ 2\ *\ +\ -$

Step	Stack Contents	Stack
Symbol read: 1		
Push 1 onto stack	1	
Symbol read: 3	1	
Push 3 onto stack	3 1	
Symbol read: 3	3	
Push 3 onto stack	3 3 1	
Symbol read: 2	3 3 1	
Push 2 onto stack	2 3 3 1	
Symbol read: *	2 3 3 1	
Pop last two items from stack (2,3)	3 1	
		Multiply 'popped' items (2,3) and store result (6) back onto stack
		Push answer onto stack
		Symbol read: +
		Pop last two items from stack (6,3) and store result (9) back onto stack
		Add 'popped' numbers (6,3) and store result (9) back onto stack
		Push answer onto stack
		Symbol read: -
		Pop last two items from stack (9,1) and store result (8) back onto stack
		Subtract second 'popped' number (9) from first (8)
		Store result

Using the stack like this is very beneficial since it allows programs to have a more efficient and lower complexity. It also allows expressions to be effectively infinitely long. As a list is traversed expressions can be pushed onto a stack until they are needed. As expressions are calculated they replace the last item on the stack.

**COPYRIGHT
PROTECTED**



Converting Reverse Polish (postfix) to infix by hand

The easiest way to convert from RPN to infix is to go through each symbol. When an operator is found, place brackets round the operator and the two arguments. For example, $5\ 4\ 6\ +\ -$ would become $5\ (4\ 6\ +)\ -$ and then $(5\ (4\ 6\ +)\ -)$. Notice how anything in brackets is converted first. Once this has been done, move the operator in each set of brackets into the middle. For example: $(5\ (4\ 6\ +)\ -)$ would become $(5\ -\ (4\ +\ 6))$. You can then remove all the brackets to get the final infix expression.

Example – Convert $10\ 4\ 2\ *\ +\ 1\ 3\ 2\ *\ +\ /\$ to infix

Step 1 – Bracket triplets starting from left

$10\ (4\ 2\ *)\ +\ 1\ 3\ 2\ *\ +\ /\$
 $(10\ (4\ 2\ *)\ +)\ 1\ 3\ 2\ *\ +\ /\$
 $(10\ (4\ 2\ *)\ +)\ 1\ (3\ 2\ *)\ +\ /\$
 $(10\ (4\ 2\ *)\ +)\ (1\ (3\ 2\ *)\ +)\ /\$
 $((10\ (4\ 2\ *)\ +)\ (1\ (3\ 2\ *)\ +))\ /\$

Step 2 – Move operators to middle of brackets

$((10\ (4\ 2\ *)\ +)\ /\ (1\ (3\ 2\ *)\ +))$
 $((10\ +\ (4\ 2\ *))\ /\ (1\ +\ (3\ 2\ *)))$
 $((10\ +\ (4\ * 2))\ /\ (1\ +\ (3\ * 2)))$

Step 3 – Remove unnecessary brackets (optional)

$(10\ +\ 4\ * 2)\ /\ (1\ +\ 3\ * 2)$

Example – Convert $((2\ 5\ *)\ 9\ +)$ to infix

Step 1 – Bracket triplets starting from left

$((2\ 5\ *)\ 9\ +)$

Step 2 – Move operators to middle of brackets

$((2\ * 5)\ 9\ +)$

Step 3 – Remove unnecessary brackets (optional)

$(2\ * 5\ + 9)$

Converting from infix to RPN by hand

To convert from infix to RPN a stack and pointer can be utilised. The first step is to identify the operators and their precedence.

High	()	Brackets
	^ or ↑	To the power
	× ÷	Multiply, divide (in computing multiply before divide)
Low	+ -	Add, subtract

Steps

1. Move pointer to the next part of the infix expression. If expression is finished, go to step 4.
2. If it is a number then write it down and go back to step 1.
3. If it is an operator follow the following:
 - 3a. If the operator is an open bracket (then push onto stack, reset the pointer to step 1.
 - 3b. If the operator is a closed bracket, pop all the operators in the stack until the open bracket is found and then discard the (from the stack.
 - 3c. Compare the current operator with the top of the stack. If it is higher precedence, push onto stack and go back to step 1.
 - 3d. Compare the current operator with the top of the stack. If it is lower precedence, pop from the stack and go back to step 3c.

COPYRIGHT
PROTECTED



Examples

Infix expression	Step	RPN	Stack	
3+2*5	1	3		Number
+2*5	3c	3	+	+ is higher than nothing
2*5	1	3 2	+	Number
*5	3c	3 2	*	* is higher than +
5	1	3 2 5	*	Number
empty	1	325*+		Pop all from stack

3+2*5 is 325*+ in RPN. When evaluated it can be seen we would work out 2*5 = 10 then 3+10 = 13.

Infix expression	Step	RPN	Stack	
(2+3)*5-(2+4)	3a		(Push (onto stack
2+3)*5-(2+4)	1	2	(Number
+3)*5-(2+4)	3c	2	+	+ is higher than nothing
3)*5-(2+4)	1	2 3	+	Number
)*5-(2+4)	3b	2 3 +	empty	Pop until (the stack
*5-(2+4)	3c	2 3 +	*	* is higher than +
5-(2+4)	1	2 3 + 5	*	Number
-(2+4)	3d	2 3 + 5 *	empty	Pop from stack
	3c	2 3 + 5 *	-	- is higher than *
(2+4)	3a	2 3 + 5 *	(Push (onto stack
2+4)	1	2 3 + 5 * 2	(number
+4)	3c	2 3 + 5 * 2	(+ is higher than *
4)	1	2 3 + 5 * 2 4	+	Number
)	3b	2 3 + 5 * 2 4 +	-	Pop until (the stack
	1	2 3 + 5 * 2 4 + -		Pop until stack empty

So (2+3)*5-(2+4) is converted to 2 3 + 5 * 2 4 + - in RPN. When evaluating it we would work out 2*3 = 6 then 6+5 = 11 then 11*2 = 22 then 22+4 = 26 then 26-6 = 20. The next stage would be 5*5 so the expression would become 5 5 * 2 4 + -. The next stage would be 5*5 so the expression becomes 25 6 - so 25 - 6 = 19. Check calculation is correct.

COPYRIGHT
PROTECTED



Task: Infix to RPN Conversion

Using the step method demonstrated on the previous page it is possible to write a program. Try to write a program that takes an infix expression as a string and converts it to RPN. For simplicity assume that all numbers are integers from 0–9 as per the example.

Questions: Reverse Polish Notation

- 1 Why is Reverse Polish notation used and why is this beneficial? (1 mark)
- 2 Convert the following into their Reverse Polish notation form. (4 marks)
 - a) $7 / 6) + 5$
 - b) $7 * 7 + (6 + 2 + 5)$
 - c) $45 / 7 + (0 - 6)$
 - d) $5 + ((1 + 2) * 4) - 3$
- 3 Convert the following to infix. (3 marks)
 - a) $a b -$
 - b) $g h b + -$
 - c) $b m / g h / +$

**COPYRIGHT
PROTECTED**



3.4 SEARCHING ALGORITHMS

LINEAR SEARCH

The linear search algorithm is used to find a given element in a list by iterating through the list and comparing it to the condition element. Linear searches always start at the beginning of the list and continue until the condition is met or the program reaches the end of the list. We say that linear search has a time complexity of $O(n)$, which means that the time it takes to run is linearly proportional to the size of the list. This is covered in more detail in Section 4.

Pseudo	
<pre> itemPosition ← 0 itemFound ← FALSE While itemPosition < ArrayLength and itemFound = FALSE If arraySearch[itemPosition] = itemWanted then itemFound = TRUE itemPosition++ Else itemPosition++ End While If itemFound = TRUE then OUTPUT "Item found at " & itemPosition Else OUTPUT "Item not found" End if </pre>	<pre> Dim itemPosition As Integer Dim itemFound As Boolean While (itemPosition < array.Length And itemFound = FALSE) If array[itemPosition] = itemWanted Then itemFound = True itemPosition++ Else itemPosition++ End If End While If itemFound = TRUE Then Console.WriteLine("Item found at {0}", itemPosition) Else Console.WriteLine("Item not found.") End If </pre>
C#	
<pre> While ((Int itemPosition = 0 > arrayLength) && (NOT isFound)) { If (arraySearch[itemPosition] = itemWanted) { itemFound = TRUE; } Else { itemPosition++; } } If (itemFound = TRUE) { Console.WriteLine("Item found at {0}", itemPosition); } Else { Console.WriteLine("Item not found."); } </pre>	<pre> Var itemFound: boolean itemPosition: integer itemWanted: integer begin itemPosition:=0; itemFound:=false while (itemPosition < array.Length) and (itemFound=false) begin if arraySearch[itemPosition] = itemWanted then itemFound = true else itemPosition++ end; if itemFound then writeln("Item found at {0}", itemPosition) else writeln("Item not found.") end. end. </pre>
Python	
<pre> itemFound = False itemPosition = 0 while itemPosition < len(arrayLength) and not found: if array[itemPosition] == itemWanted: found = True break itemPosition = itemPosition + 1 if itemFound == True: print('item found at %', itemPosition) </pre>	

INSPECTION COPY

COPYRIGHT
PROTECTED



BINARY SEARCH

A binary search is a method for searching a **sorted** set of elements. The concept is as follows:

- If the element being searched for is greater than the middle element, the middle element is discarded.
- If it is less than the middle element, then everything greater than the middle element is discarded.
- If it is the same, then the element has been found!

This process is repeated until either the element has been found or the set of elements can be further divided. Binary searches have a run time complexity $O(\log n)$.

For example, say you wanted to search for the letter S in the following set, sorted into alphabetical order:

You would start by checking against the middle item, which is F. S is to the right of F in the alphabet, so every letter to the left of F can be discarded. The letter F itself can also be discarded. This leaves the following set:

The middle element is now Y. S is to the left of Y in the alphabet, so every letter to the right of the Y can be discarded, as can the Y. The set now only consists of one letter:

The middle of the set is, now, the only item in the set and it is the letter S. The element searched for, so it has only taken three steps to complete the search. A linear search would have taken 10 steps.

When searching, the index of the middle of the set is calculated by adding the left and right indices together and dividing by 2. Any remainder is discarded. For example, if the leftmost was 1 and the rightmost was 19, then adding the two together would be 20. Dividing 20 by 2 gives 10. Therefore the index of the middle would be 10. Another way of saying this is that the index should be rounded down to the nearest integer.

The major disadvantage of the binary search algorithm is its requirement that the data has been sorted. This is a problem because sorting the array may take more time than the search! In addition, binary searches are not very efficient when performed on linked lists (as it requires going through the whole linked list), although for sets of items that are sorted and can be searched you would generally use a binary tree instead of a linked list.

Pseudo	
<pre> PROCEDURE BinarySearch (Array, ItemWanted) Left ← 1 Right ← ArraySize Middle ← 0 ItemFound ← FALSE WHILE (Left ≤ Right) AND (NOT ItemFound) Middle ← round_down((Left + Right)/2) IF Array[Middle] = ItemWanted THEN ItemFound ← TRUE END IF IF Array[Middle] > ItemWanted THEN Right ← Middle - 1 END IF IF Array[Middle] < ItemWanted THEN Left ← Middle + 1 END IF END WHILE END PROC </pre>	<pre> Procedure BinarySearch Dim bot As Integer Dim mid As Integer Dim top As Integer Dim isFound As Boolean While (bot <= top) mid = (bot + top) \ 2 If arrayName[mid] = ItemWanted Then isFound = True Break Else If arrayName[mid] > ItemWanted Then top = arrayName[mid] - 1 Else If arrayName[mid] < ItemWanted Then bot = arrayName[mid] + 1 End If End while End Procedure </pre>

COPYRIGHT
PROTECTED



C#

```

Static Void BinarySearch(arrayName, itemWanted)
{
    Int bot = 1;
    Int mid = 0;
    Int top = arrayName.Length ();
    Bool isFound = False;

    While(bot <= top) && (Not isFound)
    {
        Mid = rnd((bot + top)/2);
        If (arrayName[mid] = itemWanted)
        {
            isFound = TRUE;
            Break;
        }
        Else If (arrayName[mid] > itemWanted)
        {
            mid = mid - 1;
        }
        Else If (arrayName[mid] < itemWanted)
        {
            Bottom = Mid + 1
        }
    }
}

```

```

Procedure BinarySearch(arrayName, itemWanted)
var
    bot:Integer = 1;
    mid:Integer = 0;
    top:Integer;
    isFound:Boolean;
begin
    top:= length(arrayName);
    While (bot <= top) && (Not isFound)
    begin
        Mid := trunc((bot + top)/2);
        If arrayName[mid] = itemWanted then
            begin
                isFound := TRUE;
                Break;
            end
        Else if arrayName[mid] > itemWanted then
            Top := arrayName.Length - 1;
        Else
            Bot := mid + 1;
    end;
end;

```

Python

```

def BinarySearch(array, itemWanted):
    bottom = 0
    top = len(array)-1
    isFound = False
    while bottom <= top and (not isFound):
        mid = (bottom + top) // 2
        if array[mid] == itemWanted:
            isFound = True
            return "Found at " + str(mid)
        elif array[mid] > itemWanted:
            top = mid - 1
        elif array[mid] < itemWanted:
            bottom = mid + 1
    return "Not found"

```

Binary search using recursion

A binary search also lends itself nicely to a recursive technique as it effectively halves the size of the list. This is shown in the pseudocode below.

```

PROCEDURE Bsearch(min, max, itemRequired);
    mid = (min + max) DIV 2
    IF max < min THEN
        PRINT "not found"
    ELSEIF list[mid] == itemRequired THEN
        PRINT "FOUND IN SLOT " + mid
    ELSEIF list[mid] > itemRequired THEN
        Bsearch(min, mid-1, itemRequired)
    ELSE
        Bsearch(mid+1, max, itemRequired)

```

'max' and 'min' are index values that let us keep a track of what part of the list we are currently searching. 'itemRequired' is the value that we are looking for in the list.

COPYRIGHT
PROTECTED

This works by looking at the element in the middle of the list, and seeing whether it is the value we are looking for. If not, it compares the value found against the value that needs to be found. If the value is greater than the value we are looking for, it looks at the left half of the list. If the value is less than the value we are looking for, it looks at the right half of the list. This process is repeated until the value is found or the list is empty. The value +1s and -1s are used to disregard the midpoint, as if we have got this far, it is not the value that we want!

The terminators are the lines checking whether max is less than min, as if the list is empty (can you see why?), and the lines checking whether the element on the left is less than the element on the right. If this evaluates to true, we have found our value!

Time complexity of the binary search algorithm

The binary search algorithm is a big improvement on the linear search algorithm. The number of comparisons a binary search requires is the smallest integer value x that satisfies the equation:

$$x \geq \log_2 n$$

This is because the worst-case time complexity of the binary search algorithm is $O(\log_2 n)$. It is especially efficient when used on large amounts of data because the number of comparisons required is much smaller than for the linear search algorithm.

The table below shows the number of comparisons required for the linear search algorithm:

Number of elements	Worst case for <i>linear</i> search	Worst case for <i>binary</i> search
100	100	
1,000	1,000	
1,000,000	1,000,000	

Questions: Searching Algorithms

- 1 Consider this array:

4	3	1	6	8	9	2
---	---	---	---	---	---	---

- a) Why can't a binary search be performed on this array in its current state? (1 mark)
- b) Fix the array so that a binary search can be performed. How many comparisons would it take to find the number 2 (2 marks)
- 2 Search the following array for the letter 'R' using the binary search algorithm. The first values are given. (2 marks)

Index	2	3	4	5	6	7
A	C	E	J	L	O	Q

Left: 1

Middle: 5

Right: 10

**COPYRIGHT
PROTECTED**



3.5 SORTING ALGORITHMS

BUBBLE SORT

This sort runs up the list comparing each pair of data items. If they are the wrong way round, they are swapped over. In this way the largest item *bubbles* to the top. Next time round the second from the top item. At the same time the smallest item is moved down. The worst case run time complexity of $O(n^2)$. How this is calculated will be described in more detail later. In other words, if you have n elements to sort, in the worst case you will have to swap $n^2/2$ times.

Suppose this is the original data.

9	17	6	88	28	91	12	3	95
---	----	---	----	----	----	----	---	----

The following table shows the data changing places through one pass.

No change	9	17	6	88	28	91	12	3	95
Swap	9	6	17	88	28	91	12	3	95
No change	9	6	17	88	28	91	12	3	95
Swap	9	6	17	28	88	91	12	3	95
No change	9	6	17	28	88	91	12	3	95
Swap	9	6	17	28	88	12	91	3	95
Swap	9	6	17	28	88	12	3	91	95
No change	9	6	17	28	88	12	3	91	95

When to stop?

The code continues to make passes through the table. It can either do it eight times (even if the data is already sorted), or it can make a note of whether any swaps were made during the last pass, and stop if it wasn't. To do this, a variable `swapped` is used, which is set to true each time a swap is made, but is set to false before each pass through the data. If at the end of a pass through the data then no swaps were made, which means the sort can finish.

Pseudo	
<pre> Swapped ← True Iterations ← 0 i ← 0 temp ← 0 iterations ← 1 While iterations <= array.Length - 1 AND swapped Swapped ← False For i = 0 To array.Length - 1 - iterations If array(i) > array(i + 1) Temp = Array (i) Array(i) = array(i + 1) Array(i + 1) = temp Swapped = true End If Next Iterations++ End While </pre>	<pre> Private Sub BubbleSort Dim swapped As Boolean Dim iterations As Integer Dim i As Integer Dim temp As Integer swapped = True iterations = 1 While ((iterations < array.Length) AND swapped) swapped = False For i = 0 To array.Length - iterations - 1 If array(i) > array(i + 1) temp = array(i) array(i) = array(i + 1) array(i + 1) = temp swapped = True End If Next iterations = iterations + 1 End While End Sub </pre>

INSPECTION COPY

COPYRIGHT
PROTECTED



C#	
<pre> Static Void BubbleSort (int [] array) { Bool swapped; int iterations; int temp; swapped = true; iterations = array.Length - 1; while ((iterations > 0) && swapped) { Swapped = false; For (int I = 0; I < iterations; I++) { If (array[I] > array[I + 1]) { Temp = array[I]; array[I] = array[I + 1]; array[I + 1] = Temp; } } iterations--; } } </pre>	<pre> Repeat swapped:=false; pointer:=1; while pointer<array.Length begin if lst[pointer] > lst[pointer+1] begin temp:=lst[pointer]; lst[pointer] = lst[pointer+1]; lst[pointer+1] = temp; swapped:=true; end; pointer:=pointer+1; end; until swapped=false; </pre>

Python
<pre> def bubbleSort(myList): for passnum in range(len(myList)-1): for i in range(passnum): if myList[i]>myList[i+1]: temp = myList[i] myList[i] = myList[i+1] myList[i+1] = temp myList = [81,21,16,5,44,87,50,43,20] bubbleSort(myList) print(myList) </pre>

MERGE SORT

The merge sort is a recursive technique which follows the basic 'divide and conquer' principle. It divides the data into smaller lists until each contains just a single element. It then sorts each of them individually, then merge the elements back together so that the original data is sorted. This is accomplished in two steps: the merge and the *sort*.

The following pseudocode breaks down the two steps further:

```

Function MergeSort(listInput)
    Var left, right, result
    If length(listInput) ≤ 1 Then # single element is already sorted
        Return listInput
    Middle ← length (listInput) / 2 # find middle by dividing by 2
    For I ← 1 to middle - 1
        Add I to left # first resulting group becomes 'left'
    End for
    For I ← middle to length(listInput)
        Add I to right # second resulting group becomes 'right'
    Left ← MergeSort (left) # recursive call; input is 'left'
    Right ← MergeSort (right) # Recursive call; input is 'right'
    result ← Merge (Left, Right)
    Return result

```

**COPYRIGHT
PROTECTED**



```

...
If last (left) ≤ first(right)    # binary operator
    Append right to left
    Return left
End if
Result ← MyMerge(left, right)    # merge function in
Return result

Function MyMerge (left, right)
    Var result
    While length(left) > 0 and length(right) > 0    # while
        If first(left) ≤ first (right)    # comparison step
            Append first(left) to result
            Remove(first(left))
        Else
            Append first(right) to result
            Remove(first(right))
        End If    # end comparison step
        If length(left) > 0
            Append left to result
        Else if length(right) > 0
            Append right to result
        End If
    Return result

```

The merge sort has the average complexity of $O(n \log n)$ and the best case complexity is $O(n)$ if the array is pre-sorted. It is $O(n \log n)$ because we split at most $\log n$ times, and each time we merge the two halves.

3.6 DIJKSTRA'S SHORTEST PATH ALGORITHM

In 1956, computer scientist Edsger Dijkstra published his research findings in a paper describing an algorithm which had applications for finding the shortest path between all pairs of vertices in a weighted graph. The results are outputted by removing them from a queue of possible traversals.

```

FUNCTION DijkstraAlg (Graph, Source)
    # initialisation steps
    Dist (Source) ← 0    # set the start node distance to 0
    FOR EACH v in Graph    # for all vertices in the graph
        IF v != source Then    # if the vertex isn't the source node
            Dist(v) ← ∞    # set the distances of that node to infinity
            Prev[v] ← NULL    # set previous node to NULL (or -1)
        END IF
    END FOR
    Sum ← V + Q    # Sum the value of all nodes currently in the queue
    END FOR    # results in all vertices that aren't the source node

    # the main loop body
    While Q != empty    # while the queue isn't empty
        U ← v in Q with min(Dist(v))    # u is the vertex with the minimum distance
        Remove U from Q    # remove that node from the queue
        FOR EACH v neighbouring u    # for all neighbouring nodes
            altNode ← dist(u) + length(u,v)    # altNode compares
            IF altNode < dist(v) Then    # if altNode is less than dist(v)
                dist(v) ← altNode    # then altNode because the
                prev[v] ← u    # previous v is set to u
            END IF
        END FOR
    END WHILE
END FUNCTION


Return dist[ ], prev [ ]    # returns the distance and list of previous nodes
END FUNCTION

```

**COPYRIGHT
PROTECTED**



Take a look at the following worked example which will find the shortest path between points a and z by searching all vertices.




If we want to traverse between the nodes a and z we can use the pseudocode. At the start of our initialisation stage our source is set to node a and has the value of 0 (that is, the distance from a to a is 0). All other nodes are said to have the value of 'infinity'. We can then start to search for the shortest path.

To find the first point, we add the current distance (0) to the distances to the other nodes:

$\vec{a}b = 5$ | 5 is smaller than infinity so b is weighted as 5

$\vec{a}c = 2$ | 2 is smaller than infinity so c is weighted as 2

$\vec{a}c < \vec{a}b$ so c becomes our next node and a is removed from the queue



Now you can repeat the step again with node c as your source. For this step you add the current distance (2) to the distances to the other nodes:


$\vec{c}b = 1$ | 1 is smaller than 5 so b is reweighted as $(2 + 1 =) 3$.

$\vec{c}d = 9$ | 9 is smaller than infinity so d is weighted as $(2 + 9 =) 11$.

$\vec{c}e = 10$ | 10 is smaller than infinity so e is weighted as $(2 + 10 =) 12$.

$\vec{c}b < \vec{c}d < \vec{c}e$ so b becomes our next node and c is removed from the queue

The following node has to be d as we have no other choices and becomes weighted as 11. Node b is removed from the queue.



**COPYRIGHT
PROTECTED**




From d you could find z and end the search but you would have missed out e

$\vec{d}e = 1 \mid 1$ is smaller than infinity so e is weighted as $(9 + 1 =) 10$.

$\vec{d}z = 6 \mid 6$ is smaller than infinity so z is weighted as $(9 + 6 =) 15$.

$\vec{d}e < \vec{d}z$ so e becomes your next node and d is removed from the queue

The final step has to be to node z which has a length of 4 so it is weighted a



So, as you can see, our search algorithm would return:

a, c, b, d, e, z | distance: 14

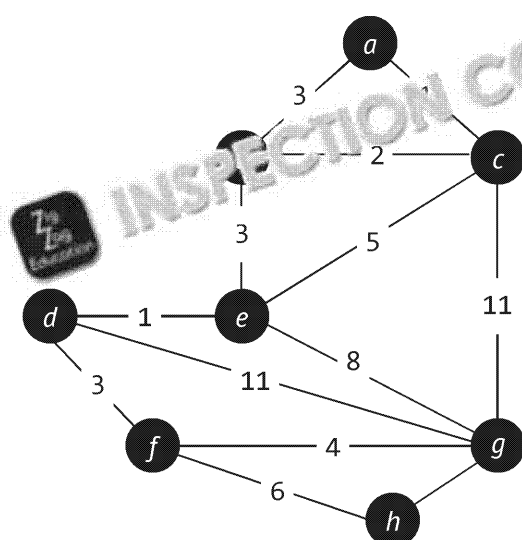
Applications

The most obvious application for the shortest path algorithm is in use with GPS devices that will find the shortest path between any two physical points on a 3D surface. The algorithm is applied automatically on the website for ease of use for the user. More specialised uses include creating a state machine to solve the shortest route to achieve a given state, or to determine the shortest time required to achieve said given state.

For example, you could solve a Rubik's cube by making each vertex a state for the cube, then solve the cube by searching for the current state of the cube and then the completed puzzle. It can also be used in networking to find the shortest path

Question: Shortest Path Algorithm

- Using the shortest path algorithm, trace the result for the graph below. Your answer should include the shortest path and the total distance.



COPYRIGHT
PROTECTED



4. Theory of Computation

The concept of computational thinking is relevant not only in the field of computer science but also in everyday life. It involves taking a problem, breaking it down into its components and abstracting the information to its essence. This section also shows the comparison of algorithms and Turing machines to show the ability to justify solutions and methods.

This section covers:

4.1 Abstraction and automation.....p1	4.4 Classification of
4.2 Regular languages.....p7	4.5 A model of comp
4.3 Context-free languages.....p14	

4.1 ABSTRACTION AND AUTOMATION

PROBLEM SOLVING

Problem solving is something the human brain is naturally adapted to do; even subconsciously while reading these words you're solving a problem of what the words mean, interpreting what you think they mean and trying to retain the information. How you apply the concept of problem solving to software development is another matter, as the scope for problems is limitless and you can't subconsciously know how to perform every action you will need to carry out.

Problem definition

The first step of problem solving is to understand the problem fully; this is called *problem definition*, and it gives you a stable foundation on which you can build your understanding of what is required. A well-defined problem will give you an understanding of the given, the resources and the end target; however, this is only the start of the problem domain.

Boundary definition

The next step is *boundary definition* which states what can and cannot be done in the problem; these act as constraints and there are a few that apply to almost all problems. Constraints are things such as time, software constraints and equipment available. All boundaries imposed on a project but do make assumptions that will improve the project. More facts by asking questions that will give you a specific answer to produce the facts to propose other constraints to the client.

Planning the solution

Stage three is *planning the solution*; in this stage you will ask yourself:

- What resources do I need?
- Are the existing resources adequate for the task?
- How will I use the resources?
- What strategies will I apply?

These all need to be addressed before you can start development. In this chapter, we will look at how it can be used to aid in planning; *decomposition* is particularly useful.

Automation

Step four is *automation* of the plans that you have generated to complete the task. It is done in a careful manner so as not to produce mistakes that may be costly to correct.

For more on software design and problem solving see Section 13.

**COPYRIGHT
PROTECTED**



FOLLOWING AND WRITING ALGORITHMS

What is an algorithm?

'An algorithm is simply a set of well-defined, step-by-step instructions that are independent of any specific programming language that can be represented in a high-level language.

This can be compared to dialling a telephone number. Knowing someone's telephone number is essential to ringing them; however, if you change the input telephone number you change the output.

You need to be aware that when an algorithm is used in a program the computer has certain restrictions on the *syntax* and the *spelling* of the algorithm and are in place. Anything that isn't a built-in instruction needs to be explained to the computer so that it can be carried out.

Hand trace simple algorithms

A *trace table* is usually used in order to investigate the flow of a simple algorithm. It usually has a column for a variable, a column for any notes, and a column for the output. It is often drawn up for any separate procedure or function that is called. This helps to find errors that could be found which by simply reading the code could be missed.

Here is an example (very simple program) with the trace table for it:

```
PROCEDURE hello()
INTEGER i
FOR i = 1 To 5
    PRINT ("Hi")
NEXT
END PROC
```

Comments	i	Output
for	1	Hi
	2	Hi
	3	Hi
	4	Hi
	5	Hi

Here is another example (a slightly longer program) with the trace table for it:

```
PROCEDURE average()
INTEGER sum, howmany, next

sum = 0
howmany = 0
READ next
WHILE next <= 100
    sum = sum + next
    howmany = howmany + 1
    READ next
WEND
PRINT "Average is " & (sum / howmany)
END PROC
```

Comments	sum
	0
read	
while	18
read	
(while)	27
read	
(while)	39
read	
(while)	45
read	
(wend)	
print	

INSPECTION COPY

COPYRIGHT
PROTECTED



Pseudocode

Pseudocode literally means 'false code' and is used to allow programmers to actually programming it fully.

Sequence	Commands are arranged and run sequentially. OUTPUT "Enter an integer: " a ← READ INPUT OUTPUT a
Assignment	Operator that assigns a value given to a variable. Unlike is replaced with '←'. x ← 5 OUTPUT x
Selection	Commands that are executed only if certain criteria are met. If-Then-Else construct. If-Else IF x ← true Then OUTPUT "Yes, it's true" ELSE OUTPUT "No, it's not true" END IF Case select Select Case letter Case letter = "C" Action ("do something") Break Case letter = "D" Action ("do another thing") break Case else Action ("do something else") End Select
Iteration	These commands are repeated in a loop until the exit condition is met. WHILE, Repeat and FOR loops. X ← 7 Y ← 5 While y > 0 Answer ← answer + (x * y) Y ← Y - 1 END WHILE

More detail of these constructs can be found in Section 1.1.

INSPECTION COPY

COPYRIGHT
PROTECTED



ABSTRACTION

Abstraction is a key component in software engineering and is something that is a generalisation of what something is, how it carries out a task and what the result is. A system that includes only the fundamental characteristics of the problem being solved represents a complex system in a way that makes it clearer and easier to understand and use. It is a way of information.

For example, a satellite navigation system may use the Dijkstra's shortest path algorithm to find the shortest path. However, the route is an abstraction of the real-life problem. The nodes in the graph are the vertices are roads. In the model the roads are weighted and shown as straight lines. In reality, roads contain bends. The actual bends in the road are irrelevant to the solution. What is important. Therefore the generalisation or abstraction is to 'ignore' the bends and treat the roads as straight weight and a straight line.

These principles are implemented by designers of computer systems in both hardware and software. In information is stored in object-oriented programming languages, abstraction is used to structure the code and objects atomic and efficient. Objects are used to represent data in a way of information to simplify the development of complex software.

Another example is a principle in mathematics called *pigeonholing* which is used as a proof for sorting using sets. For example, if you had 10 pigeons and 9 holes, how many pigeons would be in each hole? Mathematically, the formula states that if there are more pigeons than pigeonholes, then there is going to be either a pigeon without a hole or more than one pigeon in a hole.

However, with generalisation you can create the following statements:

- If x pigeons are put into y pigeonholes, and $x \leq y$, there is always a hole with one pigeon.
- If $x < y$ then there are going to be some pigeonholes with no pigeons.
- If $x > y$ then there are going to be some pigeonholes with more than one pigeon.

At its heart, generalisation and abstraction allow you to create factual statements about a system. There are entire systems that are dedicated to using just these facts in their design. These are *knowledge-based* systems where the knowledge is the culmination of the facts about a system.

INFORMATION HIDING

Information hiding, as the name suggests, is the process of hiding information about a system that is frequently observed in everyday life but is seldom noticed. When you turn on a computer, you see what files the computer is reading or caching, and this is the first type of information overload. The user of the computer isn't concerned with what files are being read as the computer turns on as expected; in fact, most users have no clue how a computer works. They notice if the process is slow or fails, but not if they are given a list of boot processes. The 'splash screen' which indicates that there is a process being undertaken and that the system is relevant to the user. A long list of filenames that they didn't know existed.

The second type of information hiding is to improve security. There is an understanding that modularisation is better; it allows the software to be worked on by separate teams. As long as the teams understand how the modules are connected, they can work on the overall task, they are not concerned with *how* the other teams are performing. They need access to their data. Variables in modules, known as *local* variables, are only used within that block of code because it isn't needed to interface with anything. This information is contained in certain blocks of code is effectively invisible and can't be accessed by other teams.

INSPECTION COPY

COPYRIGHT
PROTECTED



Did you know?!

A good example of information hiding is the Manhattan Project during World War II. The scientists, most famously, used on Hiroshima. The lead physicist knew how to build the weaponise the atomic material. The Major General at the time didn't know how to how to weaponise the atomic material for the weapons, whereas the teams that were kept separate, knew what components they were building did but didn't know why the components fitted together.

PROCEDURAL ABSTRACTION

Procedural abstraction is the act of visualising methods by abstracting actual values from subroutines. When you focus on what each task/sub-task does, you are left with a computational pattern – the programmers to focus on how the data should be handled rather than what the actual values are. Procedural abstraction is the use of *dummy* routines that alter the state of a program to continue to work towards the final solution.

For example, when creating a recursive subroutine to calculate the Fibonacci sequence, you don't worry about what the values are; you abstract them away and use variables to change and the subroutine will function. The only value you would take into

FUNCTIONAL ABSTRACTION

Functional abstraction is the abstraction of particular computational methods from the solution. You've read above that programmers will abstract actual values in order to reach function abstraction another stage is needed. Using functional abstraction, a function undergoes allows the programmer to disregard procedural blocks to prevent errors. Just like procedural abstraction, programmers can use an arbitrary value of the correct data type.

Put simply, the purpose of functional abstraction is to describe what method computation while hiding the details of how the computation is performed. In a program you can have a method that will perform an update on a class or data in state of health, but will not show how this action is performed.

DATA ABSTRACTION

Data abstraction is the next level of abstraction and is where the data type is how the action is performed. For example, if you had a class in a program that a customer you would need a routine that would retrieve the information from the class. An example of a getter method could be the user-defined method GetData() was being implemented in the class (binary tree, linked list, an array, etc.) as information in the class or class is returned the function will continue to work. An advantage is that the method can be altered without changing the code at the instance where the code is implemented. For example, in the above example as a linked list, but in an update the class code may be changed to a hash table abstraction it means that the client code wouldn't be changed and the only thing that would change is the GetData method.


**COPYRIGHT
PROTECTED**

PROBLEM ABSTRACTION

Problem abstraction relies on the premise that if you continue to remove complexity, the problem can be represented in a way that is easier to solve because at some level the problem has been solved previously.

For example, consider the idea of producing code that works out the first 30 factorials. This is everything other than the idea that you only need to work out at least two factorials. You can use recursion to calculate all factorials. Another example that illustrates abstraction is a library system. The system can be broken down into three main parts: the management of the books, the borrowing and the loan system. If you look even closer at the abstraction you can see the particular functions such as Add, Edit and Delete. This shows that books are not just data but can be saved into a file for ease of storage.

DECOMPOSITION



Decomposition is the act of breaking down a task into a set of easier identifiable subtasks. You learn from a very young age; in Mathematics, for example, when you are taught larger numbers you're taught the *divide and conquer* problem, multiplying the two halves and then adding the results.

In computer science, subtasks become even more complex and further and when all tasks are catered to, the final solution is reached.

COMPOSITION

Composition is the start of removing abstraction by beginning to form compound tasks from similar abstraction processes; for example, combining abstract procedures for borrowing and returning to form a more complex compound procedure. This has the advantage of reducing the number of tasks that form similar tasks and results in a better-structured solution.

AUTOMATION

Automation is the final step – putting all abstractions of phenomena into action to produce the final solution. This is achieved by using the abstractions you've made to design and create the algorithms, which is usually done in pseudocode to begin with.

After the planning is complete programmers will choose a suitable programming language, or it will be in the specification of the problem, to implement the pseudocode into instructions. This step includes planning what data structures are needed to fulfil the specification.

Finally, the code will need to be executed and tested thoroughly.

**COPYRIGHT
PROTECTED**



4.2 REGULAR LANGUAGES

FINITE-STATE MACHINES (FSM)

Finite state machines are a simple, intuitive way of capturing real-life events and processes. They are used by programmers to simplify and formalise the operation of programs. A lot of software is designed by using finite state machines. They can also be used as an abstract representation of the operation of Turing machines.

A finite state machine is described by:

1. A set of states
2. A start state
3. Possibly a set of final states
4. Transition function/table (dictating which inputs cause which moves)
5. Input alphabet (all the possible input events)

State transition diagrams

State transition diagrams are a way of representing finite-state machines graphically. They show the states and the transitions between them.



States – Represented by circles, states may be labelled with anything else which is appropriate as long as it is clear.



Transitions – Represented by arrows, transitions may be between different states or may loop back to the same state.

i | o or i / o


Transition Label with Output – Placed next to a transition, the input before the line and the output after the line. So in this case, the label is i | o or i / o.

i

Transition Label without Output – Placed next to a transition, the input before the line (i.e. | or /) is not needed and the output is not shown.

Here is an example of a simple state transition diagram. This finite-state machine represents a kettle. The kettle has an input alphabet of: {boiling, switch on, switch off}. It has an output alphabet of: {heater on, heater off}.

Note that at AS level you would be expected to be able to draw a FSM which requires you to draw machines with both input and output labels.



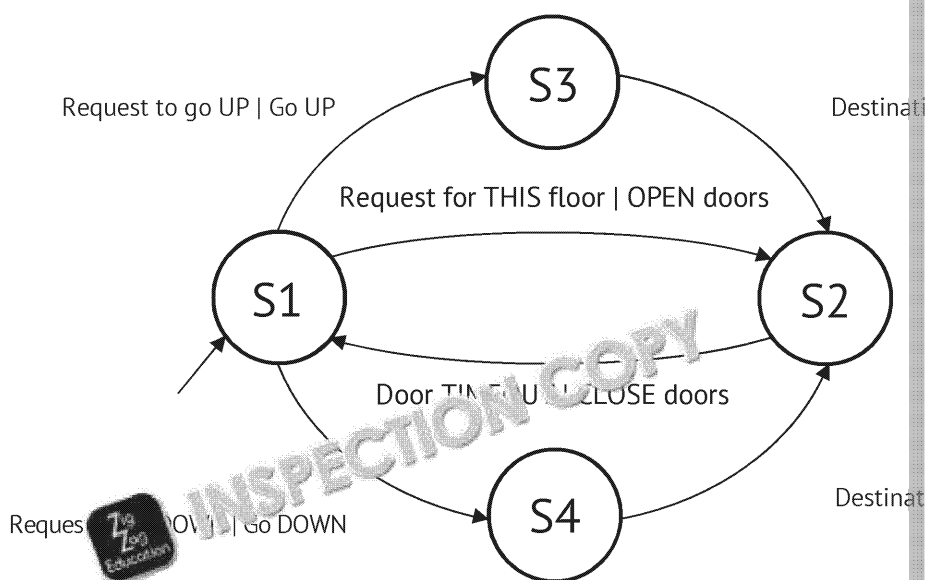
Notice that a transition can have more than one possible input/output combination. The initial state has an arrow pointing towards it.

INSPECTION COPY

COPYRIGHT
PROTECTED



Here is a slightly more complicated finite-state machine. This time it is designed to be a simple lift controller.



State transition tables

A state transition table simply maps input and state combinations to outputs. Here is an example of a transition table for the kettle:


Current State	S1	S2	S2
Input Symbol	Switch on	Switch off	Boiling
Output Symbol	Heater on	Heater off	Heater off
Next State	S2	S1	S1

Here is the transition table for the lift (note that layout isn't too important as long as the data is correct):

Current State	Input Symbol	Output Symbol(s)	Next State
S1	Request to go UP	Go UP	S3
S1	Request to go DOWN	Go DOWN	S4
S1	Request for THIS floor	OPEN doors	S2
S2	Door TIMEOUT	CLOSE doors	S1
S3	Destination REACHED	STOP and OPEN doors	S2
S4	Destination REACHED	STOP and OPEN doors	S2

Mealy machines

All of the FSMs shown so far have been Mealy machines. In Mealy machines, outputs are associated with the combination of state AND input. For a traffic light machine, it would look like this:



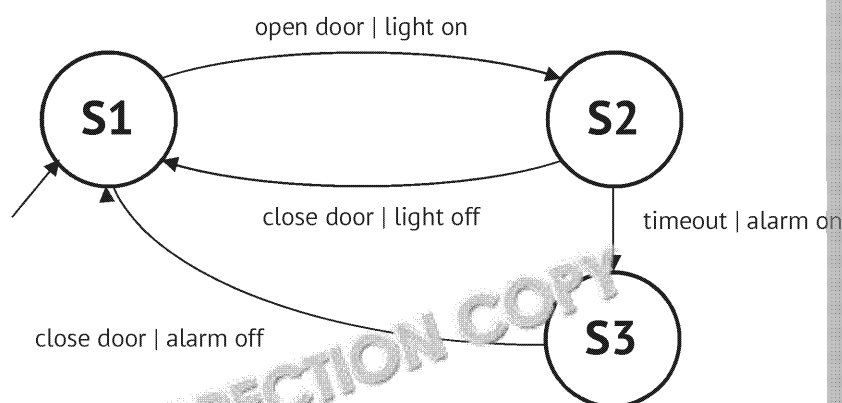
INSPECTION COPY

COPYRIGHT
PROTECTED



Questions: Finite-state Machines

- 1 This is a Mealy machine which turns on a fridge door light when the fridge is opened and sounds an alarm if it has been open too long:



- a) Give the input alphabet for this machine.
 b) Give the output alphabet for this machine.
 c) Complete a state transition table for this machine.
- 2) You are tasked with designing an automatic door system. The doors must close if the sensor has not been triggered for a certain amount of time. Design a Mealy machine which will do this task. The input and output alphabets are:
- Input alphabet = { sensor triggered, timeout }
 Output alphabet = { open doors, close doors, reset }

MATHS FOR REGULAR EXPRESSIONS

Before you can fully understand the concept and workings of regular expressions, you must come to terms with some basic concepts that you must come to terms with. This includes everything from set notation, compact representations of sets, the different types of sets and the operations on sets.

Set creation and declaration

Several programming languages support the creation of sets as built-in data types. In most languages, the default set value of any type is given the value of \emptyset meaning the empty set. For example, in Python you must use curly brackets to denote the values of the set. For example, the set of natural numbers then...

$$A = \{1, 2, 3\}$$

However, sets can also be created using *set comprehension* rules as:

$$A = \{n \mid n \in \mathbb{N} \wedge n \geq 1\}$$

In the equation above the pipe symbol '|' means 'such that', ' \in ' means 'in' and ' \mathbb{N} ' is the set of natural numbers. So the equation reads 'A is the set of numbers that are in the set of natural numbers and are greater than or equal to 1'.

As well as the objects that make the set, there are some other concepts that you need to know about sets.

- Finite sets are those sets where the values can be counted using *natural numbers*, i.e. a set of 30 objects would be counted from 1 to 30.
- Infinite sets are those where there is no end value if the range is not defined. For example, the set of *natural numbers* and the *real numbers* are both examples of infinite sets.

COPYRIGHT
PROTECTED



- *Countable infinite* sets can be counted using natural numbers in a one-to-one correspondence. You can count off all the elements in the infinite set, which, although it takes an infinite amount of time, you can index a number using its natural number match, i.e. if you have a set of integers $\{-7, \dots\}$ it is clear that it will continue forever, but you can still count them.
- The *cardinality* of a finite set is simply its size and is denoted using the symbol $|A|$. For example, the set $A = \{1, 2, 3\}$ has the cardinality 3.

You can also make new sets out of a set, where you can make bigger sets by adding elements and you can make smaller sets by removing elements from the original set. An example of a subset of the natural numbers is $\{1, 2, 3\}$.

We write this as:

$$\{1, 2, 3\} \subset \mathbb{N}$$

... to indicate that it is a proper subset, which means that the two sets are not equal.

The Cartesian product is found by 'joining' two sets together. If we take the sets $A = \{0, 1\}$ and $B = \{4, 5, 6\}$, then the Cartesian product, denoted $A \times B$, is the set

$$\begin{aligned} A \times B &= \{(a, b) \mid a \in A, b \in B\} \\ &= \{(0, 4), (0, 5), (0, 6), (1, 4), (1, 5), (1, 6), (2, 4), (2, 5), (2, 6)\} \end{aligned}$$

Set comprehension is fairly straightforward with Python – the above example can be written as:

```
A = [0,1,2]
B = [4,5,6]
AxB = [(a,b) for a in A for b in B]
print(AxB)
```

Set operators

Operator	Result
Union All of the people who have <u>either</u> blue eyes or brown hair $A \cup B = \{x \mid x \in A \vee x \in B\}$ Union (Brown hair, Blue eyes)	
Difference All the people that have brown hair but do not have blue eyes $A - B = \{x \mid x \in A \wedge x \notin B\}$ Difference (Brown hair)	
Intersection All of the people that have <u>both</u> blue eyes and brown hair $A \cap B = \{x \mid x \in A \wedge x \in B\}$ Intersection (Brown hair, Blue eyes)	
Membership Is a given individual who is a member of a certain set Membership (Fred Bloggs, Brown hair) = true	

COPYRIGHT
PROTECTED



REGULAR EXPRESSIONS

Over the course of your programming career, you'll notice that one of the most common tasks is text manipulation and pattern recognition. In *Section 1.1* we've already seen how regular expressions are used in programming, but there are limitations to where this can be used and it does not cover all possibilities (see *Section 1.2*). This is where you can use *regular expressions*; this gives you the flexibility for pattern recognition.

You can use regular expression to concisely find whether a string is formatted correctly or contains a value that you'd like to find; it can be used while reading files in text editors. The possibilities of applications are nearly endless. However, regular expressions can become complicated to read, understand and implement. *For information on how to use regular expression statements, see Section 4.2.* It is important that you know how to create and adapt to the need by combining notations.

REGULAR LANGUAGE

A regular language is one that will be accepted by finite-state machines. This is a notation to produce a set of rules to which the language will adhere to. This is where the language is comprised of its two components: its alphabet and its grammar. The alphabet is the finite set of symbols that are used and the language's syntax is ordered. *A regular language has no rules governing semantics – the meaning of the symbols.*

Regular expression notation

As regular expressions describe a set of infinite strings it's not possible to identify all valid strings using a rule, and this is where regex notation is used.

Regex notation	Meaning
a	This regular expression matches a string comprised of just the symbol 'a'.
b	This regular expression matches a string comprised of just the symbol 'b'.
ab	This regular expression matches a string comprised of the symbols 'a' and 'b'.
a*	This regular expression matches a string comprised zero or more of the symbol 'a'.
a+	This regular expression matches a string comprised one or more of the symbol 'a'.
abb?	This regular expression matches a string comprised of 'a' followed by 'b' and 'b'. There are zero or one of the symbol it follows.
a b	This regular expression matches a string comprised of the symbol 'a' or 'b'.
[a-z]	This regular expression shows a range. It includes any letter from 'a' to 'z'.
^[a-z]	This regular expression shows negation by using '^' – NOT.
[a-z] & [p-q]	This regular expression shows union and negation. The symbol 'a' to 'z' and 'p' to 'q'. It adds a condition. It reads a to z and not q.

**COPYRIGHT
PROTECTED**



Regular expression to finite-state representation

It is important that you know how to use these notations and are able to build a finite-state representation of the rules the statement is built for. These are a way of making the notation more comprehensive and understandable.


Simple statements

Simple statements where there is no function used are the most simple to represent. One symbol or string is followed immediately by another without an AND, OR, or other function. These are represented with three states and the transitions are labelled accordingly. In the example for the expression ab .



OR function (|)

When you have a regular expression it is represented by two states with two arrows between them. The second arrow you've seen previously. The transitions are labelled accordingly. In the example for the expression $a|b$. As you can see the beginning state can either transit through



The multiple function (*)

When you have a statement that states that a string will contain one or more of a symbol you use the multiple function (*). To represent this function you use a state with a transition that loops back to the state and label it accordingly. Here is an example for the statement a^* .

Piecing it all together


You can now start creating the representations for regular expressions by looking at the diagrams above to contribute to the statements diagram. You can also use these diagrams as a guide to what is valid with the regular expression rules.

Consider the expression $a|c)d^*$

We can create the finite-state representation of the expression by breaking it down into its components.

We can see by looking at it that there is a single decision (OR statement) which shows that the bracket appears first (i.e. $a|c$), directly followed by the symbol d .

Therefore, the expression can be represented by the diagram shown below.



**COPYRIGHT
PROTECTED**



You can use this diagram to verify whether strings would be valid or invalid.

Expression	Valid?
ad	Valid
accd	Invalid
cddd	Valid
cad	Invalid

Questions: Regular Languages and Expression Notation

Consider the following sets:

$$A = \{1, 4, 5, 8, 9\} \quad B = \{2, 3, 5, 6, 8, 10\}$$

1. Which set operator needs to be applied to the two sets to check what numbers are they? (2 marks)
2. What values are retrieved by the expression 'B-A'? (1 mark)
3. a) Write a regular expression that retrieves all characters apart from 'a' and 'b'. (1 mark)
b) What would the expression return for the following string? (1 mark)
the quick brown fox jumps over the lazy dog

INSPECTION COPY

INSPECTION COPY

COPYRIGHT
PROTECTED



4.3 CONTEXT-FREE LANGUAGES

BACKUS-NAUR FORM (BNF)

Backus-Naur form (BNF) is a notation which is used to create context-free grammars for natural language; these rules are what govern the syntax of a language. A natural language is naturally and is used as an everyday language with syntax rules that govern how words are constructed into phrases. These rules are based on declarations and definitions. Phrases are created; these phrases are in turn made of definitions of the constituents. These constituents can be used to construct or check whether strings/expressions are valid. For example, this would be basic English grammar for creating a list of items using ',' between items and a final item in the list. 'Egg, milk & butter' would be valid to the rules of the language, but 'milk and butter' would not.

Declarations and Definitions

The basic structure of a BNF statement consists of a *meta-component* (the thing being defined). The meta-component is enclosed in angle brackets ('<>'), and comes first in a statement, followed by a colon and a pipe symbol, which indicates that the following statements are the definition of the meta-component as follows:

```
<digit> ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
```

Similar to the pipe symbol in the pipe syntax, this expression reads that

Recursion can also be used to define an entity in terms of a previously defined entity.

```
<integer> ::= <digit> | <integer> <digit>
```

This statement defines a single digit or a digit. This is a recursive definition, as it is classed as

By utilising the recursion the integer can be any length of digits so we could define a real number as follows:

```
<real> ::= <integer>.<integer> | <integer>
```

Even though the <real> definition contains no recursion, the recursion in <integer> allows for any number of digits before the decimal point and any number after.

BNF is used to define operators, and indeed also entire expressions, functions, etc. For example, an operator and an arithmetic statement in a simple language may look like this:

```
<operator> ::= + | - | * | / | %
<statement> ::= <numericvariable> = <integer> <operator> <integer>
```

**COPYRIGHT
PROTECTED**



Limitations of regular expressions



We say a string has well-formed brackets if they match up correctly; for instance “()” is not. We want to know whether a string of brackets is well formed or not. This problem cannot be solved with just regular expressions, and this is an example where a problem cannot be solved with regex, because it does not easily deal with recursion. With BNF this is much easier, as the language of well-formed brackets can be defined as follows:

```
<lbracket> ::= (
<rbracket> ::= )
<string> ::= <string><string>|<lbracket> <string><rbracket>
```

So we can simply check whether the string can be made using these rules.

SYNTAX DIAGRAMS

The syntax of a language can also be represented through syntax diagrams. A language is defined by using entities and arrows, and any path through the diagram could be other components (the equivalent of something enclosed in angle brackets). The arrows direct a path through the entities.



Syntax diagrams can be very useful when trying to visualise complex BNF statements. They define every aspect of a programming language but the result is an intensely complex diagram that can become lost. They aren't without their drawbacks; they take up a lot of room and are not suited to being input into a computer.

COPYRIGHT
PROTECTED



Questions: Context-free Languages

1 Study the declarations below.

Sentence ::= Noun Phase Verb Phrase ;

S ::= NP VP ;

NP ::= DET Noun | Name ;

DET ::= 'The' | 'A' ;

Noun ::= 'Hippo' | 'Chair' | 'Animal' ;

Name ::= 'Keith' | 'Chris' | 'Mark' ;

VP ::= 'Sits' | 'Shouts' | 'is' ADJEC | 'is' NP, 'has'

ADJEC ::= 'long' | 'blue' | 'funny' ;

Are these sentences valid in context to the declarations above?

a) The Pen is an animal (1 mark)

b) Josh is funny (1 mark)

c) Mark is a long animal (1 mark)

d) An animal has a hippo (1 mark)



INSPECTION COPY

**COPYRIGHT
PROTECTED**



4.4 CLASSIFICATION OF ALGORITHMS

COMPARING ALGORITHMS

It is important to be able to compare algorithms as there may be multiple algorithms to solve a problem and they may vary considerably in their speed and use of memory. The speed of an algorithm is referred to as its *time complexity*. The relative amount of memory an algorithm uses is referred to as its *space complexity*. Combining these two forms of complexity gives us the algorithm's *complexity*.

The time comparison is interlinked with the space comparison because, in most cases, there are trade-offs between the advantages of both. We will often choose to improve one and as a result the other may be affected. For example, if we were finding the factors of the numbers 1 to 100, we could do this by using a loop that iterates through all integers. It is at this point that you would have to make the decision to optimise for time or space.

For space, you could make the computer store all the previous calculations it has done. This way, if a factor is required the computer will have to calculate the factor again which takes more time. To optimise it for time, you could make the computer store all previous factors for each number. For example, 64 has factors 1, 2, 4, 8, 16, 32 and 64; if we had saved previous factors we would not need to calculate all the factors of 32, so they would not need to be calculated again.

When algorithms are very simple and contain very few instructions it is easy to calculate the time complexity. For more complex algorithms, it can be difficult to work out the time complexity. Sometimes it is time-consuming and sometimes impossible when there are a number of minor optimisations. In these cases it is sufficient to calculate the complexity of the algorithm based solely on the basic operations. This operation is referred to as the *basic operation*.

MATHS FOR UNDERSTANDING BIG O NOTATION

Not all algorithms run with the same speed for all inputs. Some algorithms run faster for some inputs and more slowly with other inputs. Big O notation is the analysis of an algorithm's performance in terms of the number of basic operations it performs, considering the *worst-case scenario* and provides a notation for the *upper-bound* of the algorithm's performance. It captures the speed of an algorithm for an input that gives the worst speed of performance. An algorithm expressed in big O notation is often called the algorithm's *order of growth*.

To convert time functions into big O notation, take off the term from the function that grows the fastest. For instance, if you had the terms n and n^2 , the larger of the two is n^2 . The result obtained by n^2 is much larger than n .

Big O rules:

$$O(k) = O(1)$$

Constant times are expressed as $O(1)$.

$$O(kT) = O(T)$$

Constants inside a function are ignored.

$$O(T) + O(J) = O(T + J) = \max(O(T), O(J))$$

When adding two functions together, the bigger of the two functions is chosen.

$$O(T)O(J) = O(TJ)$$

The product of two separate functions gives the product of functions inside the big O notation.

**COPYRIGHT
PROTECTED**




Expressing complexity

Measuring the complexity of an algorithm is not as straightforward as it may seem. Measuring it is for the algorithm to be written in a programming language and timed. However, the timings generated by this method are dependent on the efficiency of the programming language. Therefore this is a crude way to measure complexity. Instead it is better to measure the speed of the algorithm based on the number of operations it requires to be carried out.

Example 1 – Two algorithms with different complexities

Consider the following problem. You are given a 3×3 grid. There are three Xs in the same row or column as another X. For example:



	1	2	3
1	X		
2			X
3		X	

The task is to locate the positions of all the Xs in the grid. So in this case the positions are (1,1) (2,3) and (3,2).


There are at least two ways to solve this problem. The first way is simply to look at each row and column of the grid to find the X. Here is some pseudocode for this algorithm:

```

FOR n = 1 to n = 3
  found ← false
  i ← 1
  WHILE found = false
    IF grid[n,i] = 'X'
      ans[n] ← i
      found ← true
    ELSE
      i ← i + 1
  END
END
PRINTLINE "The answers are:"
PRINTLINE "(1," + ans[1] + ")"
PRINTLINE "(2," + ans[2] + ")"
PRINTLINE "(3," + ans[3] + ")"

```

For the grid above, the output is:



The answers are:
 (1,1)
 (2,3)
 (3,2)

**COPYRIGHT
PROTECTED**



Assume that the instructions involving comparisons with the grid take by far the most time. These are therefore the key measure of the speed of the algorithm. This might well be the case, for example, held on a hard disk. This instruction, which is the key factor in determining the speed of the solution, is called the *basic operation*. Using this algorithm, six comparisons will always be required to find a solution, no matter how the grid is laid out. Is there an algorithm which will find a solution faster? The answer is yes. This is because not only do we know that there is only one 'X' per row, but also that there is only one 'X' per column. This means we can disregard the rows where we have already found an 'X'. One way to implement this is to use a FIFO stack to store the rows. In pseudocode, the algorithm would look like this:

```

STACK.PUSH(1)
STACK.PUSH(2)
STACK.PUSH(3)
FOR n = 1 to n = 3
  found ← false
  WHILE found = false
    i ← STACK.POP()
    grid[n,i] = 'X'
    ans[n] ← i
    found ← true
  ELSE
    STACK.PUSH(i)
END
END
PRINTLINE "The answers are: "
PRINTLINE "(1," + ans[1] + ")"
PRINTLINE "(2," + ans[2] + ")"
PRINTLINE "(3," + ans[3] + ")"

```

The number of comparisons in the above grid will now only be four! However, this algorithm is not the best. It uses a different number of comparisons for different grids. Below are two grids and the number of comparisons now required:

	1	2	3
1	X		
2		X	
3			X

3 comparisons

	1	2
1		
2		X
3	X	

6 comparisons

These two situations are called the *best-case* and the *worst-case* complexities of an algorithm. The first algorithm has a best-case complexity of three comparisons and a worst-case complexity of six comparisons.

The second algorithm has a better *time* complexity than the first on average. What is its worst-case complexity? Well, the second algorithm uses a stack, which the first doesn't. This means that the second algorithm has a worse *space* complexity than the first as it uses more memory.

Of course, in this case the difference is minor, but what if the algorithm was applied to a 100,000×100,000 grid? Also, for this example we have assumed that the stack operations take a constant amount of time compared to the comparisons; what if in reality they don't? The second algorithm would then average run a lot more slowly than the first if this were the case. Utilising memory does not necessarily speed up an algorithm.

**COPYRIGHT
PROTECTED**



Calculating the execution time of algorithms

When calculating the time complexity of an algorithm it is rarely sufficient to count the number of operations. Usually algorithms contain conditional statements and, more importantly, loops. The number of times these loops repeat is often related to the *size of the input*, denoted by n . For example, a loop that may loop through every element of an array; obviously, the larger the array, the longer it will take to execute.

When statements are simply evaluated sequentially by the computer, we say that each statement takes a constant amount of time to execute, usually denoted by c or k . Example 2 shows a sequence of simple statements expressed in this way.

Example 2 – Simple statements (i.e. the input n is constant)

Operation	Time
$A \leftarrow 3$	k
$B \leftarrow 2$	k
$C \leftarrow A + B$	k

Overall time taken is $3k$

In comparison to simple statements, each statement inside a FOR loop is executed $n + 1$ times since it requires one extra check to see whether the loop condition is satisfied. To elaborate on this point, look at example 3 and assume that n was 10. In the first iteration was taking place; j must be checked once more to see whether it is

Example 3 – Single FOR loop

Operation	Time
For $j = 1$ to n	$n + 1$
$A \leftarrow j + 1$	
End for loop	

Overall time taken is $T(n) = n + 1 + n = 2n + 1$

Nested FOR loops are slightly more complex to work out. Generally the inner loop's execution time is dictated by the number of times dictated in the outer FOR loop. Thus the execution times will be multiplied by those of the outer FOR loop. Example 4 gives an example of a nested loop to help to make this idea clearer.

Example 4 – Nested FOR loop

Operation	Time
For $j = 1$ to n	$n + 1$ times
For $i = 1$ to m	$(n + 1) * (m + 1)$ times
$A \leftarrow i + j$	$m * n$ times
End for loop	
End for loop	

Overall time taken is

Note that the constants such as 1 are not important in the long run.

For example, assume that $n = 5000$ and $m = 5,000$. There is not much difference between the expressions, $5000 + 1 + (5001) * (5001) + 5001^2$ and $5000 + (5000)^2 + 5000^2$.

Due to the small difference the constants can be ignored since they do not affect the overall time complexity.

COPYRIGHT
PROTECTED



INSPECTION COPY



Logarithmic	$O(\log n)$
-------------	-------------

- 


LIMITS OF COMPUTATION

While the development of computer hardware and software has meant that by computer have become reality, there is a limit to what is actually able to problems are where there is no algorithmic solution. For example, artificial computation; we may model behaviour using a series of algorithms but is th

Another example is paradox-type problems in which there is no decidable solution. Consider the following problem: 'There is only one window cleaner in a village. All the windows are cleaned by people in the village. Each person either cleans their own windows or let the window cleaner do it. The window cleaner is the only person who does not clean their own windows. The window cleaner cleans all and only the windows that other people do not clean themselves. Does the window cleaner clean his own windows?'

It is useful to classify problems into algorithmic or non-algorithmic to determine if they can be solved using a computer.

CLASSIFICATION OF ALGORITHMIC PROBLEMS



A problem that has an infinite set of valid inputs causes more problems as solutions are found than others may not.

One classification of algorithmic problems can be determined by its time complexity. A problem is considered solvable in polynomial time if it can be solved with a polynomial time complexity or less, i.e. $O(n^a)$ or less, is known as a polynomial-time solvable problem.

Any problem that has no polynomial time complexity or less is called *intractable*.



COMPUTABLE AND NON-COMPUTABLE PROBLEMS

Correct solutions can always be found for a solvable problem using an algorithm. A problem that is solvable, however, does not mean that they are computable in a reasonable time. A problem is solvable in less than infinite time. Unsolvability problems are those that cannot be solved in finite time, or problems that might take an infinite time to solve.

A problem which can't be solved by an algorithm

Legendre's conjecture is a seemingly simple problem, which so far has not been proven or disproven. It states that there exists at least one prime number, p , between every n^2 and $(n+1)^2$. In other words, for every n , there exists a prime p where $n^2 < p < (n+1)^2$.

One approach would be simply to loop through all the values of n and see whether there is a prime between n^2 and $(n+1)^2$. However, there is an infinite number of values of n . This means that the number of values of n is infinite, so a solution would never be found. Of course, it would be possible to find a solution for a finite number of values of n , but this would simply mean that the conjecture is true for those values of n , not that it is true.

Be wary of unsolvable problems; they are often far better disguised than this. It is a waste of time trying to come up with an answer when a workaround would be more appropriate.


Untraceable problems

Untraceable problems are problems which can be solved by a computer, but which take a reasonable amount of time for large inputs. A reasonable amount of time is generally considered to be less than or equal to $O(n)$, $O(\log n)$, or $O(n^4)$, so any algorithm that takes more than this is considered to be untraceable. Algorithms therefore have an exponential order of complexity.

An heuristic approach may be taken to help solve some untraceable problems. A heuristic is a 'guess' that can be checked in polynomial time. In other words, given a guess, the problem becomes tractable. When this is the case the problem is referred to as a heuristic problem.

The travelling salesman problem

The travelling salesman problem is a well-known problem that is difficult to solve. It is this: a salesman is trying to get through many different towns across the country. He can take that will pass through every town and go through each town only once. A graphical representation of the travelling salesman problem.



Example paths:

London \Rightarrow Plymouth \Rightarrow Bristol \Rightarrow Sheffield \Rightarrow Liverpool
London \Rightarrow Sheffield \Rightarrow Liverpool \Rightarrow Bristol \Rightarrow Plymouth
Liverpool \Rightarrow Bristol \Rightarrow Sheffield \Rightarrow London \Rightarrow Plymouth

INSPECTION COPY

COPYRIGHT
PROTECTED



The most obvious solution is to try every possible path, compare all the total distances and simply the shortest path which meets the criteria above. In computer science this is known as a *brute force* method. The issue with this approach to the problem is that it is not feasible to solve. As already explained, this means that the algorithm takes unfeasible amounts of time for large inputs (i.e. large numbers of cities). It is, however, a perfectly reasonable approach for a small number of cities (as in the diagram above), the advantage being that it finds the shortest path possible.


Exponential growth of complexity

One of the natural consequences of the growing complexity is described by the fact that the number of routes have to be taken into consideration and calculated. Taking a look at the table below, there are five cities between which all possible permutations must be calculated, using the brute force method.

Number of cities	Number of permutations
2	2
3	8
4	24
5	3,628,800
...	...
10	2,432,902,008,176,640,000

Nearest neighbour heuristic approach

A simple method of solving the travelling salesman problem for larger numbers of cities is the 'nearest neighbour'.



Select a node to start from, such as A, and choose the nearest neighbour which has not yet been visited. Repeat this process until all the nodes have been visited. If not all nodes are visited, return to the starting node and start again.

This is simplistic and does not guarantee to finish or give a good solution but it is a good starting point. The nearest neighbours from A would result in travelling the following nodes: A, G, H, J, I, F, E, D, C, B, and back to A.

**COPYRIGHT
PROTECTED**



HALTING PROBLEM

By now you'll have tried to produce some code on your own; you might even more advanced functions in your chosen language. Even if you haven't, there at some point, crashed your program. Unfortunately, it isn't possible for the going to loop infinitely. It is impossible to write a program that will detect w because you would need to know the state of the first program, which would the state of the second program.

Alan Turing (*more on on p.25*) used this as proof saying that this could go on would only halt if it halts itself, which as you know is impossible if it has ent

Another example is shown below.

```
Input x
While x > 1
    v ← x % 2 # modulo-division
    If v = 0 Then # if y = 0 then x is even
        x ← x / 2
    Else
        x ← 5x + 1
End While
```

By tracing the algorithm you can find the limits of the algorithm and see that algorithm will halt, if the input x is 5 the machine will continue to loop infinitely higher the machine will crash because it has exceeded the maximum space. We have to develop a program to test each number and the program wouldn't know infinite loop, which is why a second program would be required.

Another way of looking at it is to assume that you have a debugging program. A program along with its input and is meant to halt the program if it enters infinite loop. A program which checks whether or not the other program has stopped is called the *checker* program and its input and waits for the program to return true if the program has failed. Suppose that the program goes in an infinite loop and never returns for the response keeps on waiting until it obtains one. This will never happen and the program never terminated and so the *checker* is also never terminated since it is waiting

Questions: Classification of Algorithms

- 1 What is the run time complexity of the following algorithm? (count the number of operations)

```
FOR i=1 to n
    s = s + n
END FOR
FOR i=1 to n
    t = t * n
END FOR
```
- 2 What is the name given to each of the following orders of complexity?
 - a) $O(n^{10})$
 - b) $O(2n)$
 - c) $O(n)$
- 3 If an algorithm has a best- and worst-case scenario, how is the order of complexity determined?
- 4 What is the key difference that an algorithm with a worse space complexity has compared to an algorithm with a better space complexity when run as a computer program?

INSPECTION COPY

COPYRIGHT
PROTECTED



4.5 A MODEL OF COMPUTATION

TURING MACHINES

The Turing machine was the brainchild of Alan Turing during the Second World War. He worked on his research into computing and *computability*, a term he formally developed in the development of the *Turing machine*. The Turing machine is an abstract concept that can carry out any computable algorithm. Turing suggested that...

A number, sequence or algorithm is computable if, and only if, a Turing machine is capable of computing it.

Although Turing machines seem outdated in concept, it is an undeniable fact that they can do anything that is computable. This means that anything which a Turing machine can do as well, albeit over a longer period of time. It is also useful to think of modern-day computers without having to deal with the complexity of device architecture.

A Turing machine is comprised of:

- A tape divided into squares for reading and writing symbols to; the tape is infinitely long
- A head which can move left and right to read and write from the tape
- A transition table which depicts the operations performed by a Turing machine

A Turing machine is an example of a finite-state machine and therefore has a 'current' state and a 'next' state (defined within the transition table) which Turing called a state. It will also have an output and an input alphabet made of the symbols to be written and read.

The tape might look something like this, with the leftmost being the beginning of the tape.

1	#	1	1	0	1	1	□	0	#	1	□	1	#	0
---	---	---	---	---	---	---	---	---	---	---	---	---	---	---


Each square contains one symbol. For example, in the tape shown above '0' is a symbol. There is also a set of standard symbols which are used by convention to denote a blank symbol. There is also a set of standard symbols which are used by convention to denote a blank symbol. There is also a set of standard symbols which are used by convention to denote a blank symbol. There is also a set of standard symbols which are used by convention to denote a blank symbol. The □ character is a delimiting character, i.e. it is used to separate symbols. The □ character is a delimiting character, i.e. it is used to separate symbols. The □ character is a delimiting character, i.e. it is used to separate symbols.

Since Turing machines are a type of finite-state machine, the transition function is as it does for any other finite-state machine. The input in this case is the symbol read by the head. The output is the symbol to be written to the tape and the direction in which the head should move. The transition function is described in four parts (current state, symbol read, next state, tape symbol to write, direction to move). The machine is in a current state and reads a specific symbol; it proceeds to the next state and writes a symbol to the tape. These rules are written down and should be followed; if no rule is found, the Turing machine does not halt, it simply doesn't know what to do!

How does a Turing machine operate?

The input is the input to the Turing machine. The head of the Turing machine starts from the leftmost area of the input and can never go off the tape. A blank symbol is used to the right of the tape to indicate the end of the input stream. The head operates in accordance to the transition table where when an input is read, the transition table depicts what the next operation should be.

Machine controls tape movement in both directions



**COPYRIGHT
PROTECTED**



How can a Turing machine with a limited number of symbols represent other symbols?

Turing machines can use strings of symbols to encode other symbols. This is how an electronic computer where strings of binary numbers are used to represent data. Here is an example of such an encoding scheme.

Turing machines are often represented using unary, which is the simplest form of binary.

1, 11, 111, 1111, 11111


So to have two numbers on the tape, such as 5 and 8, we have:

11111 11111111

As spaces are hard to see, we represent them using \square s, so our numbers look like this:

$\square 11111 \square 11111111 \square$

When programming a Turing machine we assume the pointer starts and finishes at the start of the first number. We build a finite-state machine to perform the operations.



This Turing machine adds two unary numbers by removing the first 1 and moving the pointer back to the start of the number.

Written in transition rules this would be:

$\delta(S1, 1) = (S2, \square, \rightarrow)$
 $\delta(S2, 1) = (S2, 1, \rightarrow)$
 $\delta(S2, \square) = (S3, 1, \leftarrow)$
 $\delta(S3, 1) = (S3, 1, \leftarrow)$
 $\delta(S3, \square) = (\text{Stop}, \square, \rightarrow)$

Where the rule $\delta(S1, 1) = (S2, \square, \rightarrow)$ means if we are on state 1, and read a 1, we move the read head right finally moving to state 2.



**COPYRIGHT
PROTECTED**



Common way of using Turing machines

A common example of the use of a Turing machine is to validate an input string to describe a Turing machine that would recognise the language $x\#x\#...x\#x$. An input string could contain a zero or one followed by a # followed by a zero or

How would this Turing machine work?

Let the Turing machine start from the left-hand side and read the first symbol. If the symbol is a #, the machine would enter a reject state; if the symbol is a 0 or a 1 it would continue reading the next symbol. If the next symbol read is #, then it would continue reading; if it is something different, it would enter a reject state. The machine would continue reading the input by moving the head to the right. If the input string is 0 or a 1 is preceded by a # symbol.

Universal Turing machine

Alan Turing realised that a Turing machine could be extended so that it could first part of the tape carrying out each one on the rest of the tape as required. This would enable the actions of any other Turing machine. He called this the universal Turing machine concept is important because modern computers (of the stored program, or Von Neumann architecture) in that they read in a program that program as required. If it weren't for this revelation, hardware would have been running on it and so personal computers would not exist.

Questions: A Model of Computation

- 1 The Turing machine is an abstract concept invented by Alan Turing in 1936. Explain the concept and its significance in computer science. (10 marks)
- Why is the concept still used today? (1 mark)
 - Why is the universal Turing machine concept an important development? (2 marks)
 - Does a Turing machine calculate its next move purely based on its current state and the symbol it is reading? Explain your answer. (2 marks)

- 2 Consider a Turing machine with the following transitions (transitions: input symbol, next state, output symbol/move):

(50, 1, 50, >>)

$$(S_0, \cdot, S_1, 1)$$
$$(S1, \bullet, S2, \gg)$$

(S2, 1, S2, >>)

(S2, •, S3, <<)

 $(S3, 1, S3, \bullet)$ $(S3, \bullet, S4, \ll)$ $(S4, 1, S4, \bullet)$ $(S4, \bullet, S4, \bullet)$

- a) The Turing machine described above is in S_0 , with the read/write head on the blank square.

1	1	1	1	1	1	1	<input type="checkbox"/>	1	1	1
---	---	---	---	---	---	---	--------------------------	---	---	---

- (i) Which direction does the head first move in? (1 mark)
 - (ii) How many moves does the head make before it writes a symbol? (1 mark)
 - (ii) Does the machine ever enter S4? (1 mark)
- b) Describe what happens if the machine is in S4 and reads in a 0. (1 mark)
- c) Assuming numbers are coded in such a way that $1 = 0$, $11 = 1$, $111 = 0$, $1111 = 1$, $11111 = 0$, $111111 = 1$, what is the purpose of this Turing machine? (1 mark)


5. Data Representation

This section explores how computer systems are able to store a variety of different forms of data. This area has led to people misunderstanding the complexity with which computers seem to underpin seemingly 'complex' tasks. How information is stored accurately and securely is a key aspect of computer science.

This section covers:

5.1	Number systems.....	p1	5.4	Binary number systems.....	p4
5.2	Number bases.....	p2	5.5	Information coding.....	p5
5.3	Units of information.....	p5	5.6	Representing images.....	p6

5.1 NUMBER SYSTEMS

System	Description
Natural 	Natural numbers are the very first numbers you're taught. They are positive-only integers and belong to set \mathbb{N} . <i>Example:</i> $\mathbb{N} = \{0, 1, 2, 3...\}$
Integer	Integers, often called 'whole' numbers, are numbers that do not have a decimal component. These are inclusive of negative numbers and belong to the number set \mathbb{Z} . <i>Example:</i> $\mathbb{Z} = \{... -2, -1, 0, 1, 2...\}$
Rational	Rational numbers are those that can be expressed as a fraction. This means that all integers are rational numbers. Rational numbers belong to the set \mathbb{Q} . <i>Example:</i> $\mathbb{Q} = \{... 0.5, 1, 1.5...\}$
Irrational	Irrational numbers are those that cannot be written as a fraction. They are real numbers that cannot be expressed as a ratio of two integers. <i>Example:</i> $\pi = 3.14159...$
Real	Real numbers encompass all of the numbers sets as a set of quantities. Real numbers belong to the set \mathbb{R} .
Ordinal	Ordinal numbers are numerical values that hold the position of an element in an order. Consider a sorted array; the index of each element is an ordinal number.

Uses of number systems

It is important in a system to use the correct numbering system when performing calculations. It is also important to use the correct data type for a variable. The general rule is:

- Natural numbers for counting
- Real numbers for measurement

As real numbers take up considerably more memory due to their accuracy, it is important to use the correct data type. For this same reason you wouldn't use integers when you need high levels of accuracy (i.e. tracking bank balances).

Question: Number Systems

- 1 Why is it considered bad practice to use real numbers for counting and measuring? (2 marks)

**COPYRIGHT
PROTECTED**



5.2 NUMBER BASES

Data is inherently difficult to represent, store and display in a computer system. Words and symbols, or dots for images, or frequencies for sound, which can be converted into binary which can then, finally, be stored to memory. The representation of numbers in various forms. *For more on data representation see page 5.*

RADIX AND RADICES

Any value can be represented exactly using any base (*radix*). When writing values in *subscript* to avoid confusion. You need to be aware of, be able to use, and convert between *denary, binary and hexadecimal*.

Denary (n_{10})

The decimal value that we use every day (also known as *denary*) has a radix of 10. It is omitted. It uses *positional numbering*; it uses powers of 10 for each position.

For example, the number 947_{10} can be represented as:

$$(9 \times 10^2) + (4 \times 10^1) + (7 \times 10^0)$$

The number 1747.62_{10} can be represented as:

$$(1 \times 10^3) + (7 \times 10^2) + (4 \times 10^1) + (7 \times 10^0) + (6 \times 10^{-1}) + (2 \times 10^{-2})$$

Binary (n_2)

Binary representations are made up of groups of bits (*see p.5*) to convey a value. Similarly to denary, binary uses *positional numbering* but instead of powers of 10. When writing numbers in binary it is often useful to write down the values of each bit position.

For example, 86_{10} can be written in binary as:

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
128	64	32	16	8	4	2	1
0	1	0	1	0	1	1	0

Another example: 200_{10} can be written in binary as:

2^8	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
256	128	64	32	16	8	4	2	1
0	1	1	0	0	1	0	0	0

Hexadecimal (n_{16})

Digital systems store everything using binary values, but humans find it hard to work with. It requires 8 bits to represent in binary. In order to make it easier to work with, content in memory or make changes to a file, binary numbers are often grouped into bytes and displayed using a hexadecimal value. The hexadecimal number system is based on the numbers 0 to 9 and the letters A to F to represent the numbers 10 to 15.

Base 10	0	1	2	3	4	5	6	7	8	9
Base 16	0	1	2	3	4	5	6	7	8	9

The hexadecimal representations of the binary values would be $01010110_2 = 56_{16}$

INSPECTION COPY

COPYRIGHT
PROTECTED




CONVERSIONS BETWEEN RADICES

At some point you will need to convert integer numbers between radices and eventually become second nature. Although it may seem strange at first, the process is both simple and straightforward. There are two methods that can be used to convert integers between radices. These are the repeated *subtraction* and *division* methods. These methods work most easily applied between decimal and other radices.

Repeated subtraction

The repeated subtraction method uses the powers of a radix to reduce the number. As the powers work, you're left with the number 1 or 0. Take a look at the example of converting the number 190_{10} into base 2. You start off by finding the highest power that you can subtract from the number, then begin reducing the power by one after each subtraction. If the result is less than 0, simply use a 0. The highest number you can multiply the power by is given by the result of the subtraction. Converting into.

	190		
	128	$= 2^7 \times$	1
	62		
	$- 0$	$= 2^6 \times$	0
	62		
	$- 32$	$= 2^5 \times$	1
	30		
	$- 16$	$= 2^4 \times$	1
	14		
	$- 8$	$= 2^3 \times$	1
	6		
	$- 4$	$= 2^2 \times$	1
	2		
	$- 2$	$= 2^1 \times$	1
	0		
	$- 0$	$= 2^0 \times$	0

Read in this direction

As you can see, the continuous subtraction of the powers of 2 from the number 190, with the result of the radix you are left with, finish the conversion you read the result using the numbers you've multiplied. For example, 190_{10} is 10111110_2 .

Repeated division

Another method of converting integers between radices uses division instead of subtraction. It is easier, mechanical and more intuitive than the subtraction. It uses the idea that the remainder of the division is the same as successive subtraction by powers of the base.

In the following example we'll use 190_{10} to base 2 again.

2	190	0
2	95	1
	47	1
2	23	1
	11	1
2	5	1
	2	0
2	1	1

Read in this direction

As you can see in this example, the continuous division of the number 190 by 2, from the division that leaves you with a remainder of 0, finish the conversion between the two radices.

INSPECTION COPY

COPYRIGHT
PROTECTED



HEXADECIMAL CONVERSIONS

Denary to hexadecimal

As explained in 5.2.1, the binary representation is divided into groups of 4 bits, then labelled 8, 4, 2 and 1, respectively. You sum the values of each group up using the denary number system convention (0 to 9 and A to F). This is shown in the following example.

Example – Convert 213_{10} to hexadecimal

Step 1 – convert 213_{10} to binary

128	64	32	16	8	4	2	1
1	1	0	1	1	0	0	1

Step 2 – sum the values using the new labels

8	4	2	1	8	4	2	1
1	1	0	1	0	1	0	1
$8 + 4 + 1 = 13_{10} = D_{16}$				$4 + 1 = 5_{16}$			

Hexadecimal to denary

To convert between these two radices, simply use the hexadecimal place value of 16. For example, convert $A3_{16}$ to denary.

$$A3_{16} = (10 \times 16) + 3 = 163_{10}$$

Question: Number Bases

1 Convert the following:

- a) $26_{10} \rightarrow ?_2$ (1 mark) d) $01001001_2 \rightarrow ?_{10}$ (1 mark)
 b) $100_{10} \rightarrow ?_2 \rightarrow ?_{16}$ (1 mark) e) $188_{10} \rightarrow ?_2 \rightarrow ?_{16}$ (1 mark)
 c) $7A_{16} \rightarrow ?_2 \rightarrow ?_{10}$ (1 mark) f) $?_{16} \rightarrow 11010011_2 \rightarrow 201_{10}$ (1 mark)

2 What is the highest value that can be stored using a single byte?

INSPECTION COPY

COPYRIGHT
PROTECTED



5.3 UNITS OF INFORMATION

BITS AND BYTES

These are the two simplest units of data:

- A *bit* is the most basic unit of the data representation used in computing. It can be either 'on' or 'off' (1 or 0) on a digital circuit.
- A *byte* is a group of 8 bits with increasing value from right to left. A byte is the smallest *addressable* memory – meaning a specific byte can be retrieved accurately.

There are also two other frequently used representations – these might not be so obvious, but you should know! They are:

- *Words* are groups of bytes in a sequence. Frequently found word sizes are 16, 32, 64 and 128 bits.
- *Nibbles* (yes, really!) are groups of 4 bits. Therefore, a byte is formed from two nibbles: the first 4 are the high-order nibble and the last 4 are the low-order nibble.

You can work out how many values a bit pattern can represent by using the formula: 2^n , where n is the number of bits. For example, how many values can be represented by 3 bits? The answer is 8.

000 001 010 011 100 101 110 111

UNITS

While talking about small volumes of data, bits and bytes are fine. However, you're probably more than aware of the volumes of data used in real-world applications.

So how can you represent bigger volumes?

Denary		Binary	
<i>Kilo, k</i>	10^3	<i>Kibi, Ki</i>	2^{10}
<i>Mega, M</i>	10^6	<i>Mibi, Mi</i>	2^{20}
<i>Giga, G</i>	10^9	<i>Gibi, Gi</i>	2^{30}
<i>Tera, T</i>	10^{12}	<i>Tebi, Ti</i>	2^{40}

Historically the two naming conventions have been confused and the denary name is meant because it is easier to count in terms of denary.

For example, 1 kB = 1,000 bytes whereas 1 KiB = 1,024 bytes.

Questions: Units of Information

- 1 How many bits are there in 64 bytes? (1 mark)
- 2 How many bit patterns can be represented using 32 unsigned bits? (1 mark)
- 3 How many bytes are contained in a single tebibyte? (1 mark)

INSPECTION COPY

**COPYRIGHT
PROTECTED**



5.4 BINARY NUMBER SYSTEMS

UNSIGNED VS SIGNED BINARY

All the binary you've been shown in this course companion so far has been unsigned. It has a very large limitation; it cannot represent negative numbers, no matter how big. One way around this problem is to assign the highest-order bit to represent whether the number is positive or negative. This is called signed binary.

+45	0		0101101
-45	1		0101101

This is often called sign magnitude and it is not a perfect way of representing numbers. It allows for two values of zero: a 'positive' zero and a 'negative' zero. Also, by using the range of magnitude of a signed byte can represent the numbers 0 to 255, which is not what we want. It can represent -127 to 127.

The minimum and maximum unsigned binary values for a given number of bits, n , are:

UNSIGNED BINARY ARITHMETIC

Addition

Adding binary numbers is very similar to adding denary numbers except you carry over when you reach 1 (1 + 1 = 10). Similarly to when you carry digits when you reach the number 10, when you reach 1 you carry the 1. Here is a worked example.

Example: Using unsigned binary arithmetic, calculate the sum of the number 51 and 143.

Step 1 – Convert both values to unsigned binary

$51_{10} = 00110011$
 $143_{10} = 10001110$

Step 2 – calculate each bit, carrying any values

0	0	1	1	0	0	1	1		+
1	0	0	0	1	1	1	0		
<hr/>									
1	1	0	0	0	0	0	1		
1	1	1	1	1	1	1			

Step 3 – Convert the answer to denary

$11000001_2 = 193_{10}$

INSPECTION COPY

COPYRIGHT
PROTECTED



Multiplication

If you want to multiply 5 by 4 you could say that this is the same as adding 5 larger numbers this becomes slow. Multiplying two numbers is also very simple multiplication remembering that $0 * 0 = 0$, $0 * 1 = 0$, $1 * 0 = 0$ and $1 * 1 = 10$.

Consider the following example:

Binary	Denary
$ \begin{array}{r} 00110011 \\ \times 10001110 \\ \hline = 00000000 \\ 001100110 \\ 0011001100 \\ 00110011000 \\ 00000000000 \\ 000000000000 \\ 0000000000000 \\ 00000000000000 \\ + 000000000000000 \\ \hline 00111001001010 \end{array} $	$ \begin{array}{r} 51 \quad (1+2+16+32) \\ \times 142 \quad (2+4+8+128) \\ \hline 7242 \quad (2+8+64+1024+2048+4096) \end{array} $

If this multiplication were to have taken place in a computer with only an 8-bit bus it would have caused an overflow error (there wouldn't have been enough bits to represent the result). In binary arithmetic by hand it is possible to simply add the additional bits required.

SIGNED BINARY USING TWO'S COMPLEMENT

Subtraction

As shown earlier, you can represent a negative value by assigning a sign bit. However, representing a negative number without losing magnitude. If you invert all the bits of a number, the resulting binary form behaves like a negative number. This is called the number's one's complement. It is inherently difficult to use because it results in an offset of -1 . Adding 1 to a number in its one's complement, results in the number's two's complement, which is much easier to use. Follow the steps below:

Example: Using two's complement, calculate the sum of $24 - 18$.

Step 1 – Calculate the two's complement of 18 to get -18

0 0 0 1 0 0 1 0	18
1 1 1 0 1 1 0 1	One's Complement
0 0 0 0 0 0 0 1	+ 1
1 1 1 0 1 1 1 0	Two's Complement

Step 2 – Calculate the sum using the two's complement of 18

$$\begin{array}{r}
 00011000 \\
 + 11101110 \\
 \hline
 00000110 \\
 111111
 \end{array}$$

Step 3 – Convert the answer into denary

$$0000110_2 = 6_{10}$$

**COPYRIGHT
PROTECTED**



The main advantage of two's complement, however, is in calculations. This is the method.

Imagine that a number is like a counter with so many digits. If you imagine a counter going backwards, as you reach 0 and go back it would revert to 9999 for -1, 9998 for -2, etc. to 'subtract by adding'.

9997 = -3
9998 = -2
9999 = -1
0000 = 0
0001 = 1
0002 = 2

Etc.

So in base 10 we can work out 7 + (-3) by adding: $0007 + 9997 = 10004$

The additional 1 at the beginning cannot exist in a four-digit way and is called over to give the answer 0004. Two's complement does the same as the denary so equivalent to '9999' on our counter.

7 in binary is 00000111 (8 bit)

To work out -3 in binary:

+3 is 00000011

One's complement: 11111100

+ 1 becomes: 11111101

We can now add the two numbers together:

```

00000111
+ 11111101
100000100
  
```

First digit is overflow and ignored answer is 00000100 = 4

Questions: Binary Number Systems

- Using unsigned binary, complete the following:
 - $16_{10} + 44_{10} = ?_2$ (1 mark)
 - $7_{10} * 8_{10} = ?_2$ (1 mark)
 - $74_{10} + 63_{10} = ?_2$ (1 mark)
 - $9_{10} * 10_{10} = ?_2$ (1 mark)
- Convert the following into their two's complement form and complete the following:
 - $-5_{10} = ?_2$ (1 mark)
 - $-48_{10} = ?_2$ (1 mark)
- Can -166 be represented using a single byte? (1 mark)

**COPYRIGHT
PROTECTED**



NUMBERS WITH A FRACTIONAL PART

An integer is a whole number: 1, 2, 3, 4, 5, 6, etc. Decimals are numbers with a fractional part, e.g. 0.25 or 4.2039. You call the second part of these numbers the fractional part; 0.25 is the fractional part of 4.25. In computer terms this creates a problem with representation because our binary system only deals with whole numbers. In order to get round this problem a fixed- or floating-point representation is used. The parts in front of and after the decimal place are distinguishable.

There are two ways you can solve this problem:

1. *Fixed-point decimals*: allocate one set of bits for the integer part and another for the fractional part (e.g. 2.25 would be represented as 0010.0100)
2. *Floating-point numbers*: put the number into standard form and then allocate a certain number of bits for the integer part, and the rest for the fractional part (e.g. $2.4 = 2.4 \times 10^0$ which could be represented as 2.4×10^0)

Fixed-point binary

A decimal number, e.g. 11.6875, is made up of the integer part (11) and the fractional part (0.6875). In a fixed-point representation, you allocate a certain number of bits for the integer part, and the rest for the fractional part.

integer part . fractional part

The first digit of the fractional part represents $\frac{1}{2}$, the second digit represents $\frac{1}{4}$, the third digit represents $\frac{1}{8}$, etc. So for example the binary number 1011.1011 can be displayed as:

8	4	2	1	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{1}{16}$
1	0	1	1	1	0	1	1

$$\begin{aligned}
 &= 8 + 2 + 1 + \frac{1}{2} + \frac{1}{8} + \frac{1}{16} \\
 &= 11 \frac{11}{16} \\
 &= 11.6875 \text{ to 4 decimal places}
 \end{aligned}$$

If more bits in a memory word are assigned to the fractional part, greater precision is achieved. However, fewer bits are then available for the integer part and this reduces the range of numbers that can be represented. Increasing the proportion of a word given to the integer part increases the range of numbers that can be represented, but reduces the possible level of precision.

To convert the fractional part from decimal to binary you use a similar procedure to converting integers, but multiplying by 2 rather than dividing:

$$\begin{aligned}
 &0.671875 \\
 &\equiv 0.671875 \times 2 = 1.34375, \quad 0.34375 \times 2 = 0.6875, \quad 0.6875 \times 2 = 1.375, \\
 &\quad 0.375 \times 2 = 0.75, \quad 0.75 \times 2 = 1.5, \quad 0.5 \times 2 = 1
 \end{aligned}$$

Writing out the 1s and 0s in order, this gives 101011. In this case there is an integer part, so the final result is 1011.1011. This is the same as the decimal number 11.6875.

However, suppose you took a random six-digit decimal: 0.328774. Using four binary digits you get 0.3125. Even using eight binary digits you get 0.328125 which is 0.328125.

INSPECTION COPY

COPYRIGHT
PROTECTED



Negative floating-point numbers can be represented using the two's complement of the least significant bit, NOT actually 1. So for example:

8	4	2	1	1/2	1/4	1/8	1/16
0	0	1	1	1	0	1	1

represents +3.6875. To convert to -3.6875 we would perform the operation on the bits not there. So:

+3.6875 = 0011.1011

One's complement = 1100.0100

+1 to the least significant bit: 1100.0101

So -3.6875 = 1100.0101

Notice that if two's complement is being used, the maximum positive number is very quickly run out of range unless more bits are used.

Floating-point binary

Floating-point numbers are a representation of rational numbers in a binary format. The main difference between floating-point numbers rather than fixed-point numbers is that the range of numbers that can be represented with a set number of bits is far larger. Floating-point numbers achieve this by using a format similar to that used in scientific notation. Rather than have an exponent base of 10, however, most floating-point standards have an exponent base of 2. Floating-point numbers are divided into two parts: a *mantissa* and an *exponent*. They also need to have a way to represent the sign, which can be done with a single bit to represent the sign (a sign bit) or using the two's complement. An example of a real decimal number with the sign, mantissa and exponent identifier is shown below.



As the exponent base is defined by the standard being used, it is not necessarily 10. The same is true of the binary point.

In general the mantissa is specified as a fixed-point binary number. The binary point is located between the most significant bit and the least significant bit.

Converting a decimal real number to a floating-point representation

No matter which standard is used, the steps involved in converting from a decimal number to a floating-point number are very similar:

1. Convert the number from decimal to binary
2. Change binary format to the mantissa and exponent format
3. Perform two's complement conversion if numbers are to be negative
4. Normalise the mantissa and adjust the exponent so that the number is in the range of the standard


INSPECTION COPY

COPYRIGHT
PROTECTED




Two's complement floating-point numbers

One way of representing real numbers is to use the two's complement standard. This allows the representation of both negative numbers (by using numbers smaller than 1 (by using a negative exponent). Here is an example might be represented using 12 bits, 8 for the mantissa and 4 for the exponent.



The number -1.25 represented in this way would be:



The sign bit is relatively straightforward; a 0 means the number is positive, and a 1 means it is negative. The exponent is in a format called excess -127 ; this means that the number is represented as a normal binary number, but then 127 should be subtracted from it. For example, 00000001 would be -126 and so on. This is a simple way to represent numbers. Unless all the bits in the mantissa are 0, the mantissa is assumed to begin with a 1. For example, $0101101...$ is equivalent to $1.0101101...$. In effect the mantissa therefore has 24 bits rather than the 23 bits that are actually used.

The standard also includes ways to represent $+0$, -0 , $+\infty$ and $-\infty$. This means that if a number is larger than can be stored, it can be stored as ∞ .

ROUNDING ERRORS

Examples of rounding errors are shown in the next two sections.

ABSOLUTE AND RELATIVE ERRORS

There are two ways of classifying errors: absolute errors and relative errors. Absolute errors are simply the difference between the number wanted (desired value) and the computer representation of that number:

$$\text{Absolute error} = \text{desired value} - \text{computer representation}$$

For example:

$$0.492 = 1710.1992 - 1710.15$$

If the computer couldn't represent 0.1992 but could represent 0.15 the computer representation of 0.492 would be 1710.15.

The relative error is the ratio of the absolute error over the number wanted (desired value):

$$\text{Relative error} = \text{absolute error} / \text{desired value}$$

For example:

$$2.87 \times 10^{-5} = 0.0492 / 1710.1992 = 0.00002876992$$

**COPYRIGHT
PROTECTED**



RANGE AND PRECISION

The limited size of memory within a computer used for floating-point numbers limits the number of numbers which can be represented. Precision is a measure of how close the stored value within the computer is to the actual value. Say, for example, a calculation results in a value of 0.123456789 and this had to be stored as 0.123458 due to a combination of rounding error and the limitations of the floating-point system. The loss of precision would be 0.000001211.

Storing irrational numbers

Not all fractions can be represented by a finite number of bits in base 2 form
cannot be represented exactly.

1/10 is represented by 0.1 in decimal. However, in binary it becomes an infinite number, 0.0001100110011001100110011... In the same way, 1/3 is represented by 0.010101... in binary. Computers represent numbers using a finite number of bits and therefore cannot represent these fractions exactly. Fractions such as 1/3 and 1/10 will be rounded to a certain *significant bits* depending on the size of storage. The precision of the representation increases with more bits that are used.

CANCELLATION ERRORS

Cancellation errors occur when two floating-point numbers are subtracted to result is unchanged from the larger of the two numbers.

One way in which cancellation errors occur is when two numbers of completely opposite sign are subtracted. An example of this type of cancellation is subtracting 0.9999999999 from 1. The result obtained would be 1 instead of a theoretical 0.

One more way in which cancellation can occur is when you subtract two numbers. For example, if you subtract 1.000 from 1.000 the result would be a very small number instead of 0 due to the approximated representation of 1.000. This can cause division, since $1 / (1.000 - 1.000)$ would result in infinity.

Cancellation errors are usually avoided by rearranging the equation in such a way that the terms are added and then subtracted.

NORMALISATION OF FLOATING-POINT FORM

Normalisation is all about maximising the precision of the number within the mantissa. The rule is that leading digits in the mantissa should be removed, and the exponent increased or decreased as possible. Normalisation is extremely important; take the table below as an example.

Mantissa	Exponent	Normal
0000000011001	010000	No
0.11001	000100	Yes

Both of these examples represent the same two's complement number, 12.5 requires a much larger mantissa than the second to keep the same precision. limit on the size of the mantissa, say 8 bits as above, the normalised number rounding the number that has not been normalised would mean that it would


Normalising binary numbers which are not using the two's complement standard. All leading zeros in the mantissa should be removed and the exponent adjusted to be the same. As a guide, the number of places the binary point moves left by is the number the exponent should be increased by and vice versa.

Positive two's complement numbers can be treated in the same way as other important exception – there should always remain a single leading zero. Other than this, it is a completely different negative one! To normalise negative values, all leading zeros should be removed, with the exception of one.

So, in summary, to normalise positive two's complement numbers all the leading zeros should be removed, with the exception of a single zero. For example, take the following number in two's complement:

0	0	0	0	0	1	0	0	1	0
---	---	---	---	---	---	---	---	---	---

To normalise this number, remove all the leading zeros, except one and move the binary point immediately after the new most significant bit.



The exponent is the number of places the binary point has been moved to the right. In this case the binary point was moved four places.


0	1	0	0	1	0	1	0
---	---	---	---	---	---	---	---

Mantissa

The process for normalising negative two's complement numbers is very similar. For positive numbers, the leading zeros should be removed. For negative numbers, the leading ones should be removed. Take the following number as an example:

1	1	1	1	1	1	0	0	1	1
---	---	---	---	---	---	---	---	---	---

Again, the leading ones should be removed, except one, and the binary point moved immediately after the most significant bit:



Since the binary point has been moved five places right, the exponent will be 5. The final normalised floating-point number is:

1	0	0	1	1	1	1	0
---	---	---	---	---	---	---	---

Mantissa

**COPYRIGHT
PROTECTED**



UNDERFLOW AND OVERFLOW

Underflow

Underflow occurs when you are using very small numbers and the computer can store. For example, if the smallest number and you attempted the following sum:

$$1/128 \times 1/128$$

The computer cannot store the resulting number and so the computer flagging that there has been a loss in precision.

Overflow

Overflow is when the result of a numeric calculation becomes too large for the space reserved for numbers. Usually some indication will be given to indicate overflow – some machines will have a variable called an overflow flag. After a calculation if overflow has occurred. An example of trying to compute the factorial of 100 on your calculator. An overflow was shown in the binary subtraction example on page 13.

5.5 INFORMATION CODING SYSTEMS

Communication between computer systems used to pose a large problem in the institutes used different standards so computer systems would almost be talking different languages. Coding schemes allow the written characters you use every day to be converted into a form that can be manipulated, displayed and transmitted via a computer system. The main two that are used are ASCII and Unicode. While representing symbols and characters in a computer system is fairly straightforward, transferring the data can represent a bit of a problem as one computer might use a different character set to another computer. This is where standardised character sets are introduced to uniquely identify a character from a decimal value to allow communication of data without the need to define an entire character set for each file.

ASCII AND UNICODE

ASCII character set

ASCII (The American Standard Code for Information Interchange) is a 7-bit character set standard in 1963 and was widely used to represent symbols and characters. Remember that the numbers are the denary values; 65 in the ASCII character set is the value for the character 'A'.

0	NUL	16	DLE	32	SPC	48	0	64	@	80
1	SOH	17	DC1	33	!	49	1	65	A	81
2	STX	18	DC2	34	"	50	2	66	B	82
3	ETX	19	DC3	35	#	51	3	67	C	83
4	EOT	20	DC4	36	\$	52	4	68	D	84
5	ENQ	21	NAK	37	%	53	5	69	E	85
6	ACK	22	SYN	38	&	54	6	70	F	86
7	BEL	23	ETB	39	'	55	7	71	G	87
8	BS	24	CAN	40	(56	8	72	H	88
9	HT	25	EM	41)	57	9	73	I	89
10	LF	26	SUB	42	*	58	:	74	J	90
11	VT	27	ESC	43	+	59	;	75	K	91
12	FF	28	FS	44	,	60	<	76	L	92
13	CR	29	GS	45	-	61	=	77	M	93
14	SO	30	RS	46	.	62	>	78	N	94
15	SI	31	US	47	/	63	?	79	O	95

Values 8, 9, 10, and 13 convert to backspace, tab, linefeed and carriage return character representation but may affect the visual display of text. The remaining of the first 31 characters are used within the computer system itself, i.e. pointing to other characters.

COPYRIGHT
PROTECTED



Unicode character set

Unicode has now become the industry standard, made solely to allow the use of all characters from the ASCII except that it is an industry standard that can be used to represent over 100,000 characters. As a result, the majority of the world's writing systems. Like ASCII there is a set of codes for each character. This allows for the codes to be broadcast as numbers and then converted back to characters for reading. Without the Unicode character set you wouldn't be able to use emojis.

ERROR CHECKING AND CORRECTION

While communicating over a network it is possible for bits (*packets*) of data to be lost or corrupted. If a receiving device has misinterpreted a packet of data; this might result in an error. In mathematics, to drastically reduce this chance in a technique known as error checking.

There are four techniques you're expected to describe and know the use of: *parity*, *checksums*, *handshaking* and *checksums*.

Did you know?

Handshaking is a procedure that all devices undergo before transmitting data. The handshake allows both devices to check whether the other device is ready for data transmission. During the handshake, the devices agree on a method of error checking, the rate of transfer and other criteria such as method of transmission.

How this is done is defined by the particular protocol they are using, although this is largely out of the scope of your course.

Parity checking

Parity checking is a simple technique of detecting transmission errors by using a *parity bit*. This parity bit is added to the end of a transmission and is then received in reverse order. There are two types of parity: *odd* and *even*.

Odd parity	If using odd parity, the number of 1s in the transmission is summed. If the sum is odd, the parity bit is set to 0 to ensure that there is still an odd number of 1s. If the sum is even, the parity bit is set to 1 to ensure that there is now an odd number of 1s.
Even parity	Even parity is the same as odd parity except that, after counting the number of 1s is even then a 0 is added to the end of the transmission. If the number of 1s is odd then a 1 is added to the end of the transmission, as that will mean there is now an even number of 1s.

Check digits

Check digits are used to ensure that a large number has been entered correctly. A good example is looking at the *International Standard Book Number (ISBN)* found on books. The ISBN is a 13-digit number, with the first 12 digits and a check digit created by multiplying each digit by a weight and then using a modulo-11 division of the total to check whether the result is zero.

Look at the following worked example:

Example: Consider the ISBN 1 74157 103 x.

Step 1 – Write out the ISBN and calculate the positional multiples (ignoring the check digit)

1	7	4	1	5	7
10	9	8	7	6	5
10	63	32	7	30	45

**COPYRIGHT
PROTECTED**



Step 2 – Sum the total

$$10 + 63 + 32 + 7 + 30 + 45 + 4 + 0 + 6 = 197$$

Step 3 – Use modulo-11 to calculate the remainder

$$197 \text{ MOD } 11 = 10$$

Remember: if the remainder is '0' then the check digit is 0; if the remainder is 10 then the check digit is 'A'. On these two occasions you skip step 4.

Step 4 – Deduct remainder from 11 to produce check digit

In this worked example the result is one of the exceptions. However, if the remainder is not one of the exceptions then you have to deduct the remainder from 11 – this produces the check digit.

Majority voting

Majority voting uses *repetitive transmission* to deduce where the error has occurred. It sends the same number of times to the receiving device, and applies majority voting to determine the correct value, allowing for correction to be made easily. Take a look at the following example.

If one device wants to send the letter 'K' to another it can use majority voting. The receiving device knows that the binary representation of the ASCII character 'K' is 01001011 (ASCII uses 7 bits).

The sending device would then send each bit three times, so the receiving device would receive:

111, 000, 000, 111, 000, 111, 111 (commas added for clarity)

If the receiver gets the bit pattern '000' or '111' then it assumes that the bit is correct. If it gets '011' or '100' then it deduces that there has been an error. If, for the third bit, the receiver gets '001' then it can deduce that there has been an error and applies majority voting to correct it. If there are more than 1s, the original bit was a 1. Otherwise, if there were more 0s than 1s, the original bit was a 0.

Checksum

A checksum is possibly one of the oldest validation methods in data transmission. It is used for authentication because if the checksum is invalid it means that packets have been corrupted. The checksum for a transmission is determined in one of two ways. If you have a packet of data containing a single byte of data, the byte contains 8 bits which can be combined to give a total of 256 possible combinations. Ignoring the representation for 0 leaves you with 255 possible combinations.

Steps of checksum

1. Divide the total number of bytes per packet by 256.
2. Rounding the result down if needed, multiply this number by 256 (this gives you the total number of bytes that can be represented by the checksum).
3. Deducting this number from the total number of bytes being sent leaves you with the checksum value.

The following example shows how a checksum is generated for a packet that is sending 1152 bytes of data.

1. $1152 / 256 = 4.5 \approx 4$
2. $4 * 256 = 1,024$
3. $1152 - 1024 = 128$

The checksum for the transmission would be 128.

Questions: Information

1. If you were using odd parity and the data was '010011010110', what would the checksum be?
2. What would the check digit be for the number 123456789 (2 marks)
3. Using the ASCII character 'K', what would the transmission be using the majority voting method? (2 marks)

**COPYRIGHT
PROTECTED**



5.6 REPRESENTING IMAGES, SOUND AND OTHER DATA

BIT PATTERNS, IMAGES, SOUND AND OTHER DATA

You already know that a computer can only store data in binary, but it can also represent data in binary. In binary, data can only be in one of two states – on or off, positive or negative, etc. This makes digital systems as it removes the ambiguity that computers cannot handle and makes data predictable and consistent. This means you can use bit patterns to do a plethora of things, such as representing images, transforming waveforms into audio and calculating the size of files.

ANALOG AND DIGITAL

The difference between digital and analog is in the continuity of the data. Analog data is continuous, whereas digital is discrete and has a set of fixed values.

Most analog signals are from physical devices which are constantly changing, usually as a voltage or current, but a computer would be unable to read these signals if they were not discrete values. The way the signal is converted is through sampling at regular intervals and representing its amplitude. This is called analog-to-digital conversion (ADC).



One such example of analog-to-digital conversion could be a light sensor which measures the amount of light that is falling onto it at any particular time. The voltage output is proportional to the amount of light (for example, 0 means no light, 1 full light). The conversion to digital is done by representing the value of the sensor. If only 1 bit was used then the signal would only have two possible values (00 no light, 01 1/3 light, 10 2/3 light, 11 full light); 2 bits would represent four different signals (00 no light, 01 1/3 light, 10 2/3 light, 11 full light).

Equally, a device may be controlled by a computer in a similar way; for example, a light bulb can be turned on at different brightnesses. The actual light is an analog device which can receive a continuous range of voltages to provide the light range. However, the computer could only provide a discrete range of values. The number of bits used would increase the variance in signal, and the conversion from digital to analog conversion (DAC).

The most common use for ADC and DAC is in the sampling of sound – for music, for example.

More digits means larger numbers

Unfortunately, there is a finite number of digits that a computer can store. This is determined by what a computer can represent in terms of what it can store. However, the number of bits that can be stored is equal to 2^n .

n	Bit Patterns
1	$2^1 = 2$
2	$2^2 = 4$
3	$2^3 = 8$
4	$2^4 = 16$

INSPECTION COPY

COPYRIGHT
PROTECTED



BITMAPPED GRAPHICS

Bitmapped graphics use streams of bits to store information about each individual pixel in an image. These bits encode what colour the pixel should be, and as a result bitmapped images result in large file sizes. In its simplest form, 0 encodes for white and 1 encodes for black, but the more bits used, the more colours can be used. However, that's not all that the bits code for.

The *resolution* of an image is the number of pixels that appear in each inch of the image. A higher resolution means that the quality of the image is higher; this means that the image can be zoomed and scaled without visual artefacts or distortions.

The *colour depth* of an image depends on the number of bits used to store the value of the pixel. If the bit number increases, the colour depth will increase because there are more colours available to represent what the computer is trying to display, as each combination of bits will correspond to a colour in a colour chart.

The *size* of an image determines the number of rows and columns of pixels that create the image. Size and resolution are inversely proportional to each other in the sense that when you zoom in to double the size you view the image at half the resolution.

000
000
000
001

3-bit bit

Simple bitmap file calculations

Suppose a bitmap file was 3 inches by 3 inches, had a resolution of 72 pixels per inch, and 256 colours. How big would the file be? We have all the information we need to

$$\text{Size} = 3 \times 3 \text{ inches}^2 = 9 \text{ inches}^2$$

$$\text{Number of pixels per inch}^2 = 72 \times 72 = 5,184$$

$$\text{Total number of pixels} = 5,184 \times 9 = 46,656$$

256 colours require 8 bits

$$\begin{aligned} \text{File size} &= 46,656 \times 8 \text{ bits} = 373,248 \text{ bits} \\ &= 46,656 \text{ bytes} \\ &= 45.56 \text{ kilobytes (2dp)} \\ &= 0.04 \text{ megabytes (2dp)} \end{aligned}$$

This is only a rough estimate at best. In real-world applications there will always be file headers. These headers contain metadata for the image. Metadata is the information about the image, such as the resolution, colour depth, etc. The final requirements are given by the equation:

$$\begin{aligned} \text{Storage requirements} &= \text{resolution} \times \text{colour depth} \\ \text{Resolution} &= \text{width (pixels)} \times \text{height (pixels)} \end{aligned}$$

Question: Bitmapped Graphics

- 1 Estimate the size of a bitmapped file that is 7 inches by 7 inches, contains 1 square inch and has a colour depth of 8. Give your answer in kilobytes.

INSPECTION COPY

COPYRIGHT
PROTECTED




VECTOR GRAPHICS

Vector graphics are slightly different to bitmap images. Whereas a bitmap image holds the information for each pixel in an image, a vector graphic holds the information or instructions of how to create the image.

Take the following example of a line. Instead of holding information about the pixels that make up the line, a vector graphic holds information on the start and end point of a line; any pixels that intersect that line

Bitmap Graphic



Black-and-white bitmap:

1000 0100
0010 0001



Line

VECTOR GRAPHICS VERSUS BITMAPPED GRAPHICS

The table below explains the relative pros and cons of the two image types.

Bitmap graphics	
Takes up more resource memory	Takes up considerable
Takes up more storage space	Takes up considerable
Images are less precise	Graphics are more precise
Images aren't scalable without visual artefacts	Graphics are scalable
Images use less processing power	Vectors require more processing power
Made of pixels that can't be grouped	Made of elements that can be grouped

Bitmap images can be easier to edit, because you can work with individual pixels or manipulate entire planes. This makes bitmaps the only real option for photograph editing. However, if you have a large image that you may wish to edit separately to the background, a vector image is a better option. Vector images are made of small number of colours or continuous areas of colour, making them easier to edit.

Vector graphics are useful in situations such as architecture and design where precision is more important than editing small sections of an image. Also, scaling up the image without losing quality of the image and some manipulation of vector images is easier, e.g. changing the colour of a line.

**COPYRIGHT
PROTECTED**



DIGITAL REPRESENTATION OF SOUND

Sound is an air pressure wave that causes vibrations in our eardrums which our brain interprets as sound. In order for computers to capture sound waves they need a device known as an ADC (Analog-to-Digital Converter) to convert energy from one form into another.


- The original sound wave is known as *analog* data, which is a set of continuous values.
- In contrast, *digital* data such as the wave the analog data is converted into, is made up of discrete or discontinuous quantities.

Analog-to-digital conversions

Analog-to-digital conversion (ADC) takes place by measuring the height of an analog signal at regular intervals using a transducer, i.e. a microphone, and producing a digital signal representation that can be stored in a computer.




Below is a sample of a sound wave captured by the microphone with a sample rate of 10 samples per second.



Each of the dotted lines represents a sampling point; at each of these sampling points, the height of the wave is recorded. The distance between these sampling points is known as the *sampling interval*.

Below is an example of how the analog wave could look after being converted into digital data.




INSPECTION COPY

COPYRIGHT
PROTECTED




There are actually a number of steps needed to complete a conversion from analog to digital. A computer relies heavily on a technique known as *pulse code modulation* (PCM).

1. Samples are taken from the analog at a set value of hertz but must, at a minimum, be twice the highest frequency in the analog signal. There are representations of the samples that are proportional to the original signal's value in a process known as *pulse amplitude modulation* (PAM).




2. The pulse amplitude modulation values are approximated using a fixed number of bits, n , where n is the lowest number of bits that can represent the range of the analog signal. For example, if $n = 4$ then 16 levels can be used to approximate the range of the analog signal.



3. The final step is to encode the height of each of the pulses into binary. In the example you have, $10_{10} = 1010_2$, $14_{10} = 1110_2$ and $2_{10} = 0010_2$. This results in a sequence of fixed-length binary pulses which can be stored, manipulated or transmitted and then converted back into an analog waveform.

Digital-to-analog conversion (DAC)

The reverse of the analog-to-digital conversion is what you hear when listening to any sound generated from a computer or storage device such as an MP3 player.



However, during the conversion from analog to digital it is possible that during the conversion you are left with the *staircase effect*, which is where there is a subtle loss of the original analog sound and the digital representation which can be removed by using a digital filter.

INSPECTION COPY

COPYRIGHT
PROTECTED




The Nyquist theorem

Then Nyquist theorem states that...

The sampling frequency must be at least twice as high as the highest frequency because you must have at least one data point for each half-circle of the audio wave.

What this means is that if you want a full 20 kHz audio bandwidth, your sampling rate must be at least 40 kHz. This means you need a higher sampling rate to record higher frequencies. Consider the following wave; the frequency of it is constant and hence it produces a pure tone.



The sampling rate is the number of samples taken per second. In order to reproduce a sound accurately, the sampling rate must be at least twice the highest frequency of the sound. Sounds with high frequencies are high pitched.

Sample resolution and sample file sizes

As stated earlier, to record a sound you must have an input in the form of samples. The more samples taken, the closer the reproduced sound is to the original. Similarly, if you take a longer period you can increase the 'quality' of the sound. This is the sound's *sample rate*, the number of samples taken in one second of recording. If you record in the MP3 format there is a trade-off between file size and quality. The sample rate needed to capture the human voice is much lower, but if you try playing music down a telephone line the sound quality is poor.

This leads on to *sample resolution* – the number of bits assigned to each sample. The more bits, the greater the range of frequencies that can be reproduced. This is the same as the bit depth of an image file; if you increase the number of bits used to store each pixel, the range of frequencies increases, meaning that you can reproduce a sound that is closer to the original.

File size = sample rate * sample resolution * length of recording

MUSICAL INSTRUMENT DIGITAL INTERFACE (MIDI)

MIDI is a technical standard which enables a wide range of electronic musical instruments, computers and devices to connect and communicate with one another. MIDI files take up a considerably smaller volume than other sound files because only information about the notes is stored rather than the sound itself. These data items are what instrument needs to be played, what note needs to be played, and how loudly and for how long the notes are played. These data items are conveyed in what are known as *event messages* that are used to generate the sound that is required.

The major drawback of the MIDI format is that there is very little sound quality. This is because all the sounds are generated by the sound card. Sound cards only use frequency modulation (FM) synthesis or simple wavetable synthesis to generate sounds. However, MIDI is not designed for storing high-quality sound files. Instead, it is designed to allow easy composition and editing for music.

Questions: Representing Sound

- 1 Estimate the size of the sound file of a 30-second recording with a sample rate of 44,100 Hz and a sample resolution of 16 bits. Give your answer in kilobytes. (2 marks)
- 2 Calculate the sample rate of a file 1100,000 bits in size that is 10 seconds long with a sample resolution of 10 bits. (2 marks)
- 3 In audio conversion, why can you never recreate the original analog signal?

INSPECTION COPY

COPYRIGHT
PROTECTED



DATA COMPRESSION

As our computer systems have become larger and more powerful so have the what about storing and transporting all that data? The limitation on capacity data compression. Almost all files can be compressed so that they require less to transmit and are easier to work with.


Lossy compression

Lossy compression attempts to identify seemingly redundant data and remove a common example of this is in the way MP3 files are generated. The file type that are outside of the human hearing range. This can also apply to images in may be reduced to a certain level without a visible loss in quality.

Look at the example below – the same image of an apple was saved three times at 260×260 pixels. The only difference between them is the quality percentage. The lower the percentage, the more compression is applied.



Quality @ 100%
73 KB



Quality @ 66%
14 KB

As you can see, the file has been compressed from 73 KB (100% quality) to just 14 KB (66% quality). In this example, the most suitable settings have been used, which has still reduced the file size significantly (from 73 KB to 14 KB), without

Lossless compression

As the name implies, this is a method of compressing the file without loss of text, data and programs where all the information is required. The disadvantage is that it is unable to reduce file sizes as significantly as lossy compression, it often requires the user to compress and decompress the files. A common example used on the Internet

In sound, lossless compression uses a pair of algorithms. The first is to code the waveform when the sound is not in use. The second algorithm uses the inverse algorithm to regenerate the sound when it is used for playback.

Lossless compression can use *run-length encoding* (RLE) which identifies repeating patterns and stores a code for the pattern and how many times it repeats in the file. This can significantly reduce the size of the stored file, especially in text/source code and other repeating parts.

For example, suppose we want to compress the string: "I LOOOOOOOOOOVVVVVVVVVVV". We would store each repeating O, V and E once, making our compressed string: "I 10O 10V 10E". This is over 10 characters shorter than the uncompressed version.

Another approach is to use a *dictionary-based* method which uses a type of search for the file to those stored in a data structure called a *library*. If a match is found, the string is substituted with a reference pointer to the item in the dictionary. If the string is not found, it is added to the dictionary and the reference is substituted.

INSPECTION COPY

COPYRIGHT
PROTECTED



ENCRYPTION

Encryption is the act of protecting your personal data by making it unreadable. Ciphers are special algorithms designed to convert the readable data into a jumble, rendering it unreadable. This is called *cipher text* and is completely unreadable without the key which was used to generate the cipher text in the first place.

Cipher keys are an important constituent of the ciphers working and will either be symmetric or asymmetric. Symmetric keys can be used to encrypt (go from plain text to cipher text) and decrypt (go from cipher text to plain text) the data, whereas asymmetric keys will be used for one or the other but not both. With asymmetric key it is largely impossible to decrypt the cipher text, meaning if the key is lost.

Did you know?!

When passwords are stored on a web server, they are not stored in plain text form. Instead, they are hashed using a hashing table and the hash equivalent is stored. The long-lived and widely used standard hash algorithm, which was cracked in 2012 after the algorithm was 'cracked' when it leaked 6.4 million passwords. It was cracked in a single day by operators across the world using standard algorithms used and the security associated with them.

Cryptanalysis

This is the act of trying to determine the plain-text representation from a cipher without knowing the decryption key. It takes a lot of theory and is largely based around the weaknesses of the cipher keys use to secure data. In practice it is often done using what is known as a brute force attack. If a cipher key has been narrowed down to a selection of possibilities. Although it is a slow process, it can provide a greater understanding of encryption, as well as being a useful tool in the technology industry.

Caesar ciphers


Ciphers are not a technological breakthrough. Julius Caesar was known for writing any confidential military messages in a cipher with an offset of three letters to the left. It sounds like an easy system to break, but all the syntax and semantics to the words are lost; only the spaces remained. Look at the following example.

Plaintext: We attack at dawn
Ciphertext: TB XQQXZH XQ AXTK

Without knowing that there is a Caesar cipher in use it looks almost impossible to break. It is hard to recognise that there could be one in use; you could only infer its use by recognising that the spaces are retained and that there are some repeating characters, but this isn't reliable. If you did know the Caesar cipher was being used, you could break the code by using a technique known as *brute force*.

Brute force attacks use a repetitive approach to crack the key value and decrypt the data; it is for this reason that the Caesar cipher is vulnerable to this type of attack. Our simple example is limited by the number of letters in the English language, meaning at most you will have to do 25 shifts. This is often done using a *planar table*.

The table on the right shows us that our previous example is $n = 3$.



Shifts

COPYRIGHT
PROTECTED



INSPECTION COPY

Vernam cipher

The Vernam cipher was created by Gilbert Vernam during his research into creating a truly unbreakable cipher to protect data. The cipher text is produced by the sender and receiver of the message/data to be encrypted. The keys are then combined with modulo-26, to produce a truly secret message. But what is

The keys are made of streams of randomly generated letters which have each combined with the values of the letters in the data to produce an *intermediate* applied to this intermediate to create the cipher text. Take a look at the work

Encrypting using Vernam cipher

Message	C E J O J I					
Character values	0(A)	4(E)	9(J)	14(O)	9(J)	8(I)
Key values	2(C)	11(L)	16(Q)	14(O)	7(H)	24(Y)
Message + Key	2	30	35	14	9	34
Mod addition (26)	2	4	9	14	9	8
Cipher text	C	E	J	O	J	I

Decrypting using the Vernam cipher.

Now that you've seen how to encrypt using the Vernam cipher, you need to know essentially the same process but in reverse, except where you used modular addition to encrypt, you use modular subtraction to decrypt.

Message	C E J O J I					
Character values	2(C)	4(E)	9(J)	14(O)	9(J)	8(I)
Key values	2(C)	11(L)	16(Q)	14(O)	7(H)	24(Y)
Message + Key	0	-7	-7	0	2	-16
Mod addition (26)	0	19	19	0	2	10
Cipher text	A	T	T	A	C	K

It has been mathematically proven that, when used properly, the Vernam cipher offers no information about the plaintext it represents and that, if the criteria are met, data encrypted using the cipher will be secure if, and only if:

1. The keys must be equal to or greater than the length of the original message
2. The keys must be discarded or reused
3. The keys are generated using algorithms that produce truly random numbers
4. The keys are distributed in a secure and secretive manner

Given enough material (cipher text) and time, all cryptographic algorithms, even the most secure, can be broken. This is because all other devised encryption algorithms that aren't based on computation security depend on *computation security*.

Computation security is creating security based on an exponential time algorithm. As the encryption length grows, the time taken to break the encryption increases exponentially.

**COPYRIGHT
PROTECTED**



Questions: Data Compression and Encryption

- 1 How does lossy data compression reduce the size of a file? (1 mark)
- 2 Using a left-shift Caesar cipher, decrypt the following message: (1 mark)
JVTWBALY ZJPLUJL
- 3 Why is it important that the length of a key being used for a Vernam cipher message being encrypted? (2 marks)
- 4 Is the Vernam cipher immune to brute force attacks even if the keys have been used before? (1 mark)



INSPECTION COPY



INSPECTION COPY

INSPECTION COPY

COPYRIGHT
PROTECTED



6. Computer Systems

In this section we look at the fundamental concepts of computing and how a series of electronic components have come together to form the modern-day computer that has become part of everyday life. The concepts behind this theory and the design of computer processors is vital for the modern-day computer scientist.

This section covers:

6.1 Hardware and software.....	p1	6.4 Logic gates.....	
6.2 Classification of programming languages.....	p5	6.5 Boolean algebra.....	
6.3 Types of program translator.....	p7		

6.1 HARDWARE AND SOFTWARE

RELATIONSHIP BETWEEN HARDWARE AND SOFTWARE

- The *hardware* of a system is the physical components that make up the computer system, e.g. the motherboard, power supply, hard drive, and other peripherals (plug-in devices) such as keyboards, mice, scanners, printers, are also hardware.
- The *software* of a system is a collection of procedures and rules (lines of computer code) that carry out operations on a computer's hardware. The software code may be stored on the computer's hard drive, a DVD, a USB stick or any other storage device.

It is necessary to provide input to make the computer do something. This input can come through many different means, for example your mouse, keyboard or scanner, or other devices. Less obvious input is the power button, the computer's internal battery (which has its own battery) and possibly instructions given to the computer from a user.

The software reads in the input, works out what it is that you've told it to do, and then carries out a certain specific sequence of instructions to achieve the desired result. The result is sent to an output device such as your monitor or printer.

Both hardware and software are essential for a computer to run; without the software the input given to the computer through the input devices would not be recognised – the computer won't boot up! But without the hardware the software would not be able to physically carry out the instructions.

The hardware and software in your computer are not just working when you turn it on. From the moment you press the 'on' button the two are working together to perform tasks. Even when there doesn't seem to be anything happening on the computer, they are running continual checks and maintenance to ensure that everything runs smoothly.

CLASSIFICATION OF SOFTWARE

The purpose of a computer is to store the software on its hard drive. When the computer is turned on, the software is loaded into memory. For example, you can have a computer comprised of hardware only, but without the software the computer won't start. Software can be classified into two main categories: system software and application software, depending on its purpose.

- *System software* is independent of any general-purpose package or any particular application area, but is designed to assist the computer in the efficient execution of application programs. An example of a piece of system software would be an operating system (such as Windows or Linux).
- *Application software* allows the user to achieve a specific task that the user wishes it to perform; this could be anything from a word-processing program or web browser, to the software that helps to run a nuclear power station.

INSPECTION COPY

COPYRIGHT
PROTECTED



SYSTEM SOFTWARE

System software acts as the brains behind the operation as far as managing a computer is concerned. Without it, the computer is just a collection of components that wouldn't function. System software, including the *operating system*, *library programs*, *translator software* and *utility programs*, are the software that manage the computer's hardware and provide the user with a way to interact with the computer.

Library programs

A *library program* is the generic name given to a collection of programs which are used by other programs. If you install certain modern games you will be asked if you accept the terms and conditions required for the computer to interpret the instructions required to run the game. These terms and conditions provide libraries required to run built-in services. The code in libraries is not created by the user; it is edited by independent programmers to allow them to save time having to create code from scratch. A file explorer can be used to save files instead of the programmer having to create a file from scratch.

Translator software

All software translators have a single purpose and that is to convert one programming language into another. Translators are used to convert source code created by a programmer into a form that the computer can understand. The three main types of translator are:

- *Assemblers* are used to convert a low-level language called 'Assembly language' into machine code. Assembly language used to be the only option other than coding in machine code. Words are used to represent memory locations or operations; these words are converted into machine code which is a language that the computer can use and understand, and is called machine code.
- *Compilers* check that all the lines of the program are valid to the syntax of the language and then convert the entire source code into object code. Object code is usually a low-level code which is dependent on the machine it is running on. It may be translated to machine code using assembler, linker, binders etc. The compiled code is distributed as software for a particular machine or as a set of programs for sale. The original source code would be similar with minor changes to be compiled into a particular machine format.
- *Interpreters* are similar to compilers except that they read the source code line by line, check if the statement is valid and then execute that line before moving on to the next line until the end of the source code is reached. Some programming languages are interpreted. For example, JavaScript is an interpreted language. Each platform has its own JavaScript interpreter which translates the code received into machine code. Interpreted languages are slower in operation than a compiled program but have the advantage of being able to run across platforms, i.e. to change a JavaScript code you change the program and distribute for each platform.

Utility programs

Any programs that are operated by the user to maintain the functionality of a computer are given the title *utility programs*. These are small and useful programs that are used for file management, diagnostic tools and system information tools that provide information about the system (e.g. system resource usage) about your computer.

Most utility programs that are required for safe, maintainable operation of your computer are included with your operating system software package and cover tasks such as disk defragmentation, file explorers and everything down to copy-and-paste operations. Some utility programs maintain or configure, monitor certain items on your system or enable transfer of data between different systems.

INSPECTION COPY

COPYRIGHT
PROTECTED



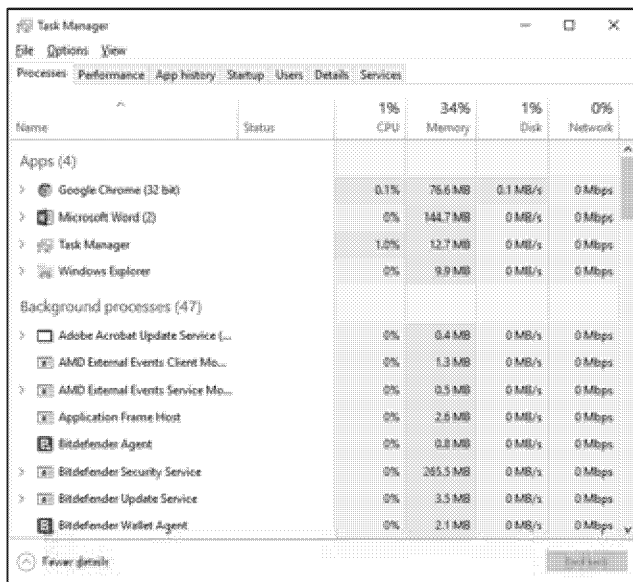
ROLE OF AN OPERATING SYSTEM

The operating system (also known as the *OS*) is the single most important and necessary piece of system software. While the computer is in use, the operating system runs in the background managing the system's resources and processes. System processes include communications with input/output devices, managing memory and managing programs – meaning the user doesn't need to worry about how the system resources are being used (unless they're doing a task which is particularly resource-intensive, e.g. 3D rendering). This acts as an interface between the user and the computer, application software and components of the system.

The operating system is not a single entity but a collection of programs that work together to supply a level of abstraction for the user. For example, a user can tell the computer to delete files from a directory and the computer interprets this and issues the corresponding instructions to carry out the task. The abstraction is achieved by creating a *virtual machine* in which the complexities of the tasks are hidden from the user.

The general role of an operating system is to provide an environment from which users can interact with the computer. This results in operating systems needing to provide the following services:

- Resource management (process management, memory management)
- User interface



Name	Status	1% CPU	34% Memory	1% Disk	0% Network
Apps (4)					
Google Chrome (32 bit)		0.1%	76.6 MB	0.1 MB/s	0 MBps
Microsoft Word (2)		0%	144.7 MB	0 MB/s	0 MBps
Task Manager		1.0%	12.7 MB	0 MB/s	0 MBps
Windows Explorer		0%	9.9 MB	0 MB/s	0 MBps
Background processes (47)					
Adobe Acrobat Update Service (...)		0%	0.4 MB	0 MB/s	0 MBps
AMD External Events Client Mo...		0%	1.3 MB	0 MB/s	0 MBps
AMD External Events Service Mo...		0%	0.3 MB	0 MB/s	0 MBps
Application Frame Host		0%	2.6 MB	0 MB/s	0 MBps
Bitdefender Agent		0%	0.8 MB	0 MB/s	0 MBps
Bitdefender Security Service		0%	265.5 MB	0 MB/s	0 MBps
Bitdefender Update Service		0%	3.5 MB	0 MB/s	0 MBps
Bitdefender Wallet Agent		0%	2.1 MB	0 MB/s	0 MBps

Process management

Process management involves the switching of program execution between processes. In operating systems, multiple processes can be running at the same time. For example, a user can play an MP3 file and watch a video. In effect, this is what allows a single processor to do many things at once.

Many modern computers have multiple processors, each of which can handle multiple processes. This means that a certain number of processes can be running simultaneously. However, in older computer systems, there was only one processor, so processes had to be switched in and out of execution at once and the user had to need to switch between tasks.

An operating system schedules programs and switches between them by managing the execution of processes. Switching between programs is done so that it appears that multiple processes are running at the same time. Since there are many processes that want to run, they have to be organised in a queue and priorities must be given to them. To be fair to the programs, all of the processes must be given a limited amount of time before then being switched in an order that affects the response time for interactive applications (delay when using mouse or keyboard). The operating system also try to finish as many processes as possible, and also balance out the system so that everyone gets a fair chance.

Users may wish to alter the priorities of processes, and this can sometimes be done using an application such as the Task Manager in Windows (shown on p.3). For example, if a user wants to edit photos, the user may want to increase the priority of the video application so that it can finish editing the photos.

INSPECTION COPY

COPYRIGHT
PROTECTED



Device management

Managing devices using an operating system is useful because the operating system can make each device accessible to programs. In order for devices to be handled appropriately by the operating system, device drivers are needed. A *device driver* is a special piece of software that controls the hardware and importantly provides an interface so that programs can use the device. In other words, device drivers give a layer of abstraction to the software which makes use of the devices.

By using this abstraction, programs can access and perform operations on the hardware via simple function calls. An example could be the following command: 'play sound from file'. Behind the scenes the file would be decoded and sent to the sound card which would then process the sound and output it to the speakers. From the programmer's point of view, the only thing that he or she is required to know is the filename.

Hardware devices are usually organised in terms of priorities. When a hardware device wants to perform a duty, an interrupt is signalled. It then decides whether or not interrupts are allowed to occur. If they are, then the operating system processes it by saving the current state of the running program, executing an interrupt handler which directs the flow of execution to the appropriate device. If a device causes an interrupt, then the operating system decides which device has the highest priority and that device to obtain the attention of the processor by executing the interrupt handler.

Memory management

Memory in a computer system is finite and used by all processes, and therefore the operating system will have to divide up and keep track of the memory that is available. Sharing a finite resource in a fair and effective way between many entities is not an easy job. In order to avoid processes running out of memory and crashing, many operating systems operate using virtual memory.

This is where the operating system moves the contents of memory to and from the hard disk intelligently – hard disks are generally much slower than RAM, and so overall performance is slowed down considerably.

Provision of a virtual machine

Abstraction is the key idea behind operating systems that makes them so popular today. By abstracting all the processes behind hardware and software, in addition to providing user-friendly interfaces, all the complicated operations behind a computer system are hidden.

A machine which is complicated in the background but easy to use due to these layers of abstraction is also classed as a virtual machine.

Questions: Hardware and Software

- Identify the software types of the following:

a) Antivirus software (1 mark)	c) C# compiler (1 mark)
b) Ubuntu (1 mark)	d) File explorer (1 mark)
- How is abstraction created between the user and the computer system?
- What is the main drawback of using an interpreter for translation? (1 mark)

**COPYRIGHT
PROTECTED**



6.2 CLASSIFICATION OF PROGRAMMING LANGUAGES

Programming languages haven't always been advanced as they are now – as we've taken baby steps to where we are now – but before you can learn about languages you must first understand the concept of *language levels*.

LOW-LEVEL LANGUAGES

Low-level programming language's defining feature is the lack of human-language abstraction between the programmer and the instructions being programmed. They are highly dependent on the design of the system and are described as *machine-level*. There are two low-level languages that you need to know about; these are *machine code* and *assembly language*.

Machine code

Machine code is the lowest coding language and is comprised entirely of 1s and 0s, representing on and off electrical states in the computer. Machine code is usually written in hexadecimal representation of the byte, because it is more concise to type, but it is still frustratingly tedious and difficult. However, because of how precisely the programs were highly dependent on what processor they were being run on.

Machine code instructions

Although when you look at machine code it just looks like a load of 1s and 0s, it's actually commands and data. Here is an example set of commands for a machine code program:

0000	X	Store accumulator value in memory location X
0001	X	Load contents of memory location X into the accumulator
0010	X	Set accumulator value to X
0011	X	Add contents of memory location X to accumulator
0100		Halt execution

The *accumulator* is a register that is used to manipulate numbers. X in this case represents a memory location (see Section 7 for more information).

Machine language example 1

Using the machine instructions above, this program adds the values in memory location 0001 and 0000, and puts the answer in 0010:

0001 0000 0011 0001 0000 0010 0100

Machine language example 2

As mentioned above, the code is shown in the hexadecimal numbering system. Here is an example of a program space on the page. The computer commands shown here (taken from machine language for an IBM mainframe computer) are instructions instructing the computer to load two numbers, compare them, and if the first is greater, move the result into the output area of memory, and set up the next instruction to be printed.

Assembly language

Assembly language is considered a very low-level language by today's standards. It was a big leap forwards in programming design because of one particular feature: it was human-readable.

Instead of the 1s and 0s used in machine language, assembly languages use abbreviations that are easy to remember. For example, an assembly language instruction to Compare, 'MP' for Multiply, STO for storing information in memory, etc. Although they still make programming a lot simpler than coding in hexadecimal, they are still quite complex.

INSPECTION COPY

COPYRIGHT
PROTECTED



Assembly language also allows for the use of rudimentary variables such as C memory address. These mnemonics are said to have a one-to-one relationship which is what allows them to be used unambiguously.

As with all languages, assembly language requires a *translator* to be read by translator software is called an *assembler program*, or simply *assembler*. It takes assembly language and converts it into machine code instructions.

Assembly language example

This example shows an assembly language extract from a programming language called A68000. Notice that the names for subroutines are given on the left, the instructions are given in the next column, and memory locations (variables) and other information are given in the third column.

An exclamation mark (!) precedes a comment. Note that this code does exactly the same as the high-level language shown in the *high-level* section.

HIGH-LEVEL LANGUAGES

High-level programming languages, as you might have guessed, contain elements of natural spoken languages. These elements are what provide abstraction and the instructions being programmed. Instead of dealing with registers, memory codes, high-level programmers deal with variables, data structures and complex pointers. The complex memory pointers are hidden from the programmer and the complexity is managed by the compiler.

Examples of high-level languages include the group of programming languages known as Imperative languages. Imperative languages are the most dominant and widespread paradigm for programming. Basic, C#, Pascal, etc. The main characteristic of these languages is that all statements are executed in sequence and the flow of the program is diverted using constructs (i.e. *iteration*) which statements are executed is crucial to the success of the code. Variable 'state' of the program which can be used to control the flow of execution.

Compilers

A translator is needed to translate the symbolic statements of a high-level language into machine language. This translator is usually a compiler. Keep in mind the following:

- A compiler converts a high-level language into machine code. As every computer has a different machine code, it follows that each different computer needs a different compiler to take advantage of the particular machine code).
- Although each compiler is different, high-level languages are more portable. A source program can be used on different machines, and the compiler will generate the correct machine code.
- Different organizations may bring out different compilers for the same language. Over time, improvements or bug fixing which result in more than one version of a compiler. Sometimes the high-level language may have differences to account for different machines. Therefore standardisation is needed to keep the portability.

**COPYRIGHT
PROTECTED**



High-level code example

Many examples of high-level code can be found throughout this resource, for Python and VB .NET. Below is an extract of C# code that is a translation of the following language example.

```
Static int sign (int i)
{
    int sign = 0;
    If (i > 0)
    {
        if (i < 0)
        {
            sign = -1;
        }
        Else
        {
            sign = 1;
        }
    }
    Return sign;
}
```

Did you know?!

Back in the day programming was a lot more complex than it is now. In the early 1990s magazines would publish code snippets for use and even algorithm puzzles in as little as 10 lines of code to trace and dissect. Suddenly Sudoku d

Questions: Programming

- 1 Using the high-levelled language example above, write a program that:
 - a) Write comments for what the program does.
 - b) State what the output would be if the number 5 was passed into the function? (5)
- 2 What are the main differences between a high-level language and machine code? (3)
- 3 If a high-level language runs on a different system? Explain your answer. (3)

6.3 TYPES OF PROGRAM TRANSLATOR

As mentioned earlier, a *translator* converts program statements written in one programming language into another programming language. The most common translators are assemblers that convert program statements written by a programmer in an appropriate language for the job (the source code) into machine code that can be executed by a computer.

Originally, programming was carried out in machine code. Programmers would type in commands using hexadecimal, which was converted to binary machine code for the computer. A major step forward was when assembly languages were written which had commands that corresponded more closely to English abbreviations (*mnemonics*).

Assemblers

An *assembler* then converted these assembly language instructions into the machine code. Early assemblers did little more than convert instruction mnemonics to their machine codes, but later assemblers did much more. These assembler tasks have to be performed by the computer to convert instructions to machine code:

- Assembly language supports the use of *macro* instructions (the equivalent of subroutines in a high-level language), they are expanded by the assembler and the code inserted in the assembly. Instructions are checked for syntactical correctness; errors are reported. The assembler is constructed to link symbolic operands and labels with their corresponding machine codes.
- Many assemblers support the use of *pseudo-operations*, or *directives*, which are not translatable into machine instructions. These involve such things as reserving space for variables, defining values of identifiers used in the program and defining where the program starts.

**COPYRIGHT
PROTECTED**



Interpreters

An *interpreter* takes one line of a high-level language source code program as an instruction for that programming language, translates the statement to machine language, and then executes the resulting machine language before translating the next program statement.

Advantages of an interpreter over a compiler are:

- The programmer can execute the program and sort out each problem even if there are many invalid lines of code later on in the program. The interpreter translates the code line by line while developing a program.
- If an error is found the source code can be corrected.

Compilers

A *compiler* checks that all the lines in the program are valid. Then, providing the whole program into machine language that can then be executed by a computer. Advantages of a compiler over an interpreter are:

- The resulting machine code runs faster than an interpreted program.
- The resulting machine code can be run on a computer that doesn't have the interpreter.
- A person given a copy of the compiled program can't see the original source code, so it's harder for people stealing the code for their own uses.
- The compiler program doesn't have to be resident in memory at the time the program is run. Only the larger programs can be run in memory.

When are translators used?

Assembler	An assembler takes assembly code, which is written in mnemonics. These are understandable by the computer and therefore can be executed directly. Assemblers are used to create machine code.
Interpreter	Translating code a line at a time and running it is very useful for languages which have both an interpreter and a compiler. It is useful for development for easy bug fixing, and the compiler when you are working to produce the executable program. Interpreters are also useful for compiling on the fly on web servers. For example, JavaScript.
Compiler	These translators are used to create an executable file of machine code. This executable file can then be run at a later time. Because the compiled code runs faster than trying to compile it each time, it can give a performance boost for computationally heavy programs.

6.4 LOGIC GATES

All modern computers work on a binary system, so it is important to have a way of formulating circuits. In logic circuits you're generally answering a 'yes' or 'no' question to formulate a final outcome. For example, a logic circuit may produce a response to the computer? by asking 'am I bored?' and 'do I have any work to do?' – if the answer to the second is no then you can play on the computer.

We may want to accept a user if he is the system manager or technician, if the user is correct, and if the user is not barred from the system. This is just one use of logic operations. In programming we apply them to *conditions* (a condition in this context is either true or false). They are often used in conjunction with relational operators, which

**COPYRIGHT
PROTECTED**



TRUTH TABLES

The following tables show the outcomes when each logical operator is applied to two conditions.

- NOT is used to negate (or change) a condition.
- AND is used to check that both of two conditions are true.
- OR is used to see whether either of two conditions is true.
- NAND (Not AND) provides the inverse of the AND operator.
- NOR (Not OR) provides the inverse of the OR operator.
- XOR (exclusive OR) is used to see if only one or the other applies.

Note: XOR is not a basic logical operator as the equivalent test can be made using AND and OR here as it is useful to see whether two conditions are different.

Here you can see the completed truth tables for all logical operators.

Condition A	Condition B	A AND B
True	True	True
True	False	False
False	True	False
False	False	False

Condition A	Condition B	A OR B
True	True	True
True	False	True
False	True	True
False	False	False

Condition A	Condition B	A NAND B
True	True	False
True	False	True
False	True	True
False	False	True

Condition A	Condition B	A XOR B
True	True	False
True	False	True
False	True	True
False	False	False

Condition A	Condition B	A XOR B
True	True	False
True	False	True
False	True	True
False	False	False



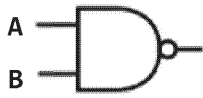
Condition A	Condition B	A NOR B
True	True	False
True	False	False
False	True	False
False	False	True

LOGIC DIAGRAMS

There is a better way of representing logical procedures, and it allows the reader to see the logic of a procedure. We call these *logic diagrams* and they make understanding a procedure easier.

Each logic gate has its own symbol and one or more inputs, and produces a single output. By putting them together it is possible to create circuits that perform a complicated logical operation.

The examples of the logic gates that have been explained using the truth tables above can be represented by logic diagrams. The examples on the left can take the values of the conditions A and B from the columns of the truth tables.

Operator	Boolean	Logic Diagram
AND	$(a \cdot b)$	
OR	$(a + b)$	
NAND	$\overline{a \cdot b}$ or $(a \cdot b)'$	

Operator	Boolean
NOR	$\overline{a + b}$
XOR	$(a \cdot b) + (\overline{a} \cdot \overline{b})$
NOT	\overline{a}

INSPECTION COPY

COPYRIGHT
PROTECTED



LOGIC CIRCUITS


Logic gate symbols can be combined to create circuits of logical operations that take an input to give an output. In a Boolean circuit the information flows from left to right. Below is an example of a simple Boolean circuit which has the following inputs:

- $A \rightarrow$ is the weather forecast sunny?
- $B \rightarrow$ is the weather forecast not windy?
- $C \rightarrow$ is the umbrella not broken?

This produces a response to the question: 'Shall I take an umbrella?'

One possible set of inputs where the circuit will produce a 'yes' is if weather is sunny (A) and neither is the umbrella broken (C) nor is there a wind forecast (B). Putting this into a Boolean expression:

$$A' \text{ AND } (B \text{ AND } C)$$



We can construct the truth table for the above Boolean circuit. There are three possible input combinations each resulting in one true (T) or false (F) output. For two inputs there would be 2^4 combinations, five inputs would mean 2^5 combinations and so on.

Input A	Input B	Input C	Output ($A' \text{ AND } (B \text{ AND } C)$)
True	True	True	False
True	True	False	False
True	False	True	False
True	False	False	False
False	True	True	True
False	True	False	False
False	False	True	False
False	False	False	False

INSPECTION COPY

COPYRIGHT
PROTECTED




Half adders and full adders

In electronics, an adder is a combination of logic gates that are used to produce the sum of two or more numbers. In modern computers adders are used in the processor and in system memory when calculating addresses for memory allocation, calculating return addresses and in many other applications.

Half adders

Half adders add two single bits that we will call *A* and *B*. It uses these two inputs to produce a sum (represented by the letter *S*) and a signal known as *carry* (*C*). The carry signal is used for producing full adders, as seen below. The example below uses an XOR gate for the sum and an AND gate for the carry. The truth table can be seen below.

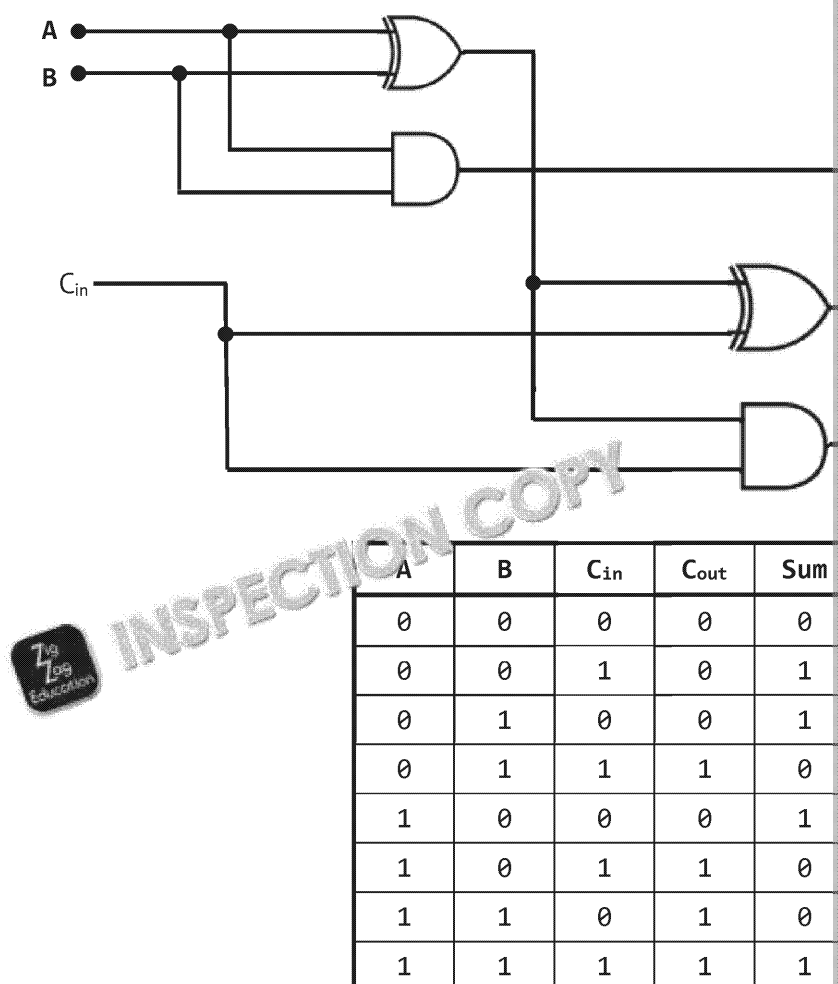


Full adders

Full adders are the combination of two or more half adders and where the half adder produces an output for sum and carry, the full adder will take the carry and use it as a carry-in for the next half adder.

Consider you're adding together two 3-bit numbers. You will need to create three full adders to create the full adder. The carry-out bit of one is when you're adding the next bit. Therefore your truth table will have five columns.

The logic gate diagram and truth table can be seen below.



COPYRIGHT
PROTECTED



Edge-triggered D-type flip-flop circuits as memory units


Flip-flop circuits can take a signal and, depending on previous states, alter the output. It sounds complicated but it is the basic principle for computer memory.

Edge-triggered D-type flip-flop circuits are used in static RAM modules. Static RAM is usually used for core system operations because, although the speed is much slower, the sizes of the modules are much smaller and they take up considerably more space.


Questions: Logic Gates

1 State the outputs of the following logic gates.


a) Where A is TRUE and B is FALSE?




b) Where A is FALSE and B is FALSE?



c) Where A is TRUE?



2 Draw a truth table for the following Boolean circuit:



3 Construct a Boolean circuit for the following: (1 mark)

$$(A+B) \cdot (C+D)$$

**COPYRIGHT
PROTECTED**




6.5 BOOLEAN ALGEBRA

In Boolean algebra there are always only two possible outcomes, true or false (represented as 0 or 1). Boolean algebra has a set of *postulates* which are basic rules for all the possible operations. In Boolean algebra we read a + as OR and \cdot as AND.


One way to visualise Boolean algebra is as simple electrical circuits with switches. A switch is on the corresponding Boolean expression is outputting a 1, and a 0 if the switch is off. Two switches on the same wire and an OR is equivalent to two switches on the same wire (or closed) then it is equivalent to a 1 and a 0 if it is off (or open).

THE BASIC RULES OF BOOLEAN ALGEBRA


1. A variable can be either 1 or 0



2. $0 + 0 = 0$ (off OR off = off)



3. $1 + 1 = 1$ (on OR on = on)




INSPECTION COPY


COPYRIGHT
PROTECTED




4. $0 \cdot 0 = 0$ (off AND off = off)




5. $1 \cdot 1 = 1$ (on AND on = on)



6. $0 \cdot 1 = 1 \cdot 0 = 0$ (off AND on = off, on AND off = off)



7. $1 + 0 = 0 + 1 = 1$ (off OR on = on, on OR off = on)



INSPECTION COPY

COPYRIGHT
PROTECTED



The Boolean algebra laws

There are 11 Boolean algebra laws, 10 of which are shown below. The eleventh is discussed in more depth in the following section. In this section A , B and C represent any Boolean expression. Remember that \bar{A} means not (A) and A , B and C can be any Boolean expression. Remember the umbrella circuit example described earlier in this section (see p.10). The laws are listed below, normally with the intention of simplifying them.

1. $A + B = B + A$
 $A \cdot B = B \cdot A$

This means that Boolean algebra is *commutative*; the order of the inputs does not matter.

2. $(A + B) + C = A + (B + C)$
 $(A \cdot B) \cdot C = A \cdot (B \cdot C)$

This means that Boolean algebra is *associative*; the order in which the sub-expressions are evaluated does not matter.

3. $A \cdot (B + C) = (A \cdot B) + (A \cdot C)$
 $A + (B \cdot C) = (A + B) \cdot (A + C)$

This law means that Boolean algebra is *distributive* so a multiplication can be distributed over an addition or vice versa. This is the equivalent of 'expanding brackets' in normal algebra.

4. $A + A = A$
 $A \cdot A = A$

This law is very different to normal algebra but is a consequence of the fact that there are only two possible outcomes in Boolean algebra so adding (or multiplying) two true values results in another true. It is known as the *identity law*.

5. $A \cdot B + A \cdot \bar{B} = A$
 $(A + B) \cdot (A + \bar{B}) = A$

Because either B or not B will be true, multiplying by each of them will result in a true outcome.

6. $A + (A \cdot B) = A$
 $A \cdot (A + B) = A$

Although this law doesn't look like it should be true, it can be proved by using the identity law: $A + (A \cdot B) = A \cdot (1 + B) = A \cdot 1 = A$. This law is known as the *redundancy law*.

7. $0 + A = A$
 $0 \cdot A = 0$

This law simply describes the effect of adding or multiplying by 0.

8. $1 + A = 1$
 $1 \cdot A = A$

If there is one true value in a Boolean expression then the outcome will be true. Adding or multiplying with a true value means that the result will always be true.

9. $A + \bar{A} = 1$
 $A \cdot \bar{A} = 0$

This law follows from the previous two, remembering that when you have a true value it will always be true (1) and the other false (0).

10. $A + (\bar{A} \cdot B) = A + B$
 $A \cdot (\bar{A} + B) = A \cdot B$

Similar to rule 9, this law holds because A and \bar{A} will always have one true value.


**COPYRIGHT
PROTECTED**




De Morgan's law

De Morgan was an English mathematician who moved to England from India and discovered that in logic it was possible to show that the following two sets of rules allows for Boolean circuits to be simplified and hence more understandable, saving space if they are used in hardware devices.

De Morgan's first law is $\overline{(A+B)} = \bar{A} \cdot \bar{B}$



De Morgan's second law is $\overline{(A \cdot B)} = \bar{A} + \bar{B}$




You can check these laws by constructing the truth tables for the circuits and comparing them to the original expression.

By using these laws it is possible to see how Boolean expressions can be simplified into a single circuit (or sub-circuit if it is part of a larger circuit) with the simpler two-gate circuit.

SIMPLIFYING A BOOLEAN EXPRESSION

Using the laws it is possible to take a typical Boolean circuit and simplify it. Below is an example of a circuit that can be simplified.




- This is the circuit for $(\bar{A} \text{ OR } \bar{B}) \text{ AND } (C \text{ OR } (C \text{ AND } D))$, which can be simplified using De Morgan's laws.
- This circuit can be simplified using the second De Morgan's law and Law 6 to simplify the $(C + (C \cdot D))$ to C .
- De Morgan's law simplifies $(A' + B')$ to $(A \cdot B)'$.

COPYRIGHT
PROTECTED



- So the expression becomes $(A.B)'C$



- This is the circuit for $\overline{(A \text{ AND } B)} \text{ AND } C$.

As you can see, this is a much simpler circuit and as result will be much simpler inside the computer.

Simplifying Boolean expressions can be performed like simplifying mathematical expressions.

Example

Simplify $A(A+B)(A+C)$

We can effectively 'multiply the brackets':

$$AA + AC + BA + BC$$

AA is equivalent to A

$$A + AC + AB + BC$$

We can see a 'factor' in the first three expressions of A :

$$A(1 + C + B) + BC$$

Using law 8 we know that 1 or anything is 1:

$$A.1 + BC$$

Using law 9 we know that 1 and A is A :

$$A + BC$$

As can be seen, the simplified expression would be less complicated to construct.

Questions: Boolean Algebra

1 Use the laws of Boolean algebra to simplify the following: (5 marks)

a) $A + (A \cdot B)$

b) $A \cdot B + A \cdot C$

c) $A + A \cdot B + (A + A \cdot B)$

d) $A + \overline{B}$

e) $\overline{(A + B)} + B$

**COPYRIGHT
PROTECTED**



7. Computer Organisation and Architecture

Computer architecture is an extension of the logic shown in the previous sections. Knowledge of how the machine actually handles information is a vital part of understanding how a computer operates.

This section covers:

7.1 Internal hardware components.....	p1	7.3 Structure and role of the operating system.....	p1
7.2 The stored program concept.....	p5	7.4 External hardware components.....	p1

7.1 INTERNAL HARDWARE COMPONENTS

INTERNAL STRUCTURE AT A GLANCE


So far you have already covered what is required in order for components of a computer to allow users to use the computer to complete a task (see Section 6.1). Even though the list of components found in the average machine has increased over time, the basic components remain the same. When you open up your computer, inside you will find the following:

Motherboard

Everything in the computer is in some way connected to the motherboard, be it via slots, wires, connectors or readouts; everything must connect to the central processing unit (CPU) board. The motherboard acts as a central interface for all components; not all components connect to the board in the same way, though. How a device is connected depends on the bandwidth of the device, how often it is used and what the device is used for.

For example, the *Peripheral Component Interconnect Express (PCI-E)* lanes that graphics cards are not the same as those used to connect hard drives.

Hard disk drive (HDD) and solid-state drives (SSD)



For a long time, HDDs have been the most common way for a computer to store its programs and data. They are large and slow. A computer thinks there is more than one hard drive if there are multiple operating systems to be installed simultaneously. HDDs are large in volume but are not that fast compared to SSDs.

Nowadays, SSDs are more being used more increasingly, either alongside an HDD or as the primary storage. Despite being more expensive than HDDs, their write/speeds are much faster, allowing the system to boot in just seconds.

Power supply

The *power supply unit (PSU)* provides and regulates the supply of power to the computer. It provides power to all other components. Modern PSUs allow the computer to control when it turns on and off, saving energy and for the user to define how they want the power to be delivered to the system.

Power supplies come in varying sizes – not just in terms of their physical size, but in the number of watts they can support; the standard power supply on an average computer is roughly between 200 and 500 watts but there are high-tier PSUs available that can deliver up to 1,500 watts!

INSPECTION COPY

COPYRIGHT
PROTECTED



How the pieces are connected

As mentioned previously, all the components are connecting to the motherboard. The method of communicating along the wires connecting components to the motherboard may be restrictive, i.e. may limit the speed at which the computer can operate.

The hardware components on the motherboard are also connected together using printed circuitry called *buses* (see p.3 for information).

The connections are vital for different parts of the computer to work together for instance loading a program code from the hard disk into RAM when an app is launched.

A CLOSER LOOK AT INTERNAL STRUCTURE

Processor

The Central Processing Unit (CPU) is the brains of the computer and is the component used for all everyday applications. It manipulates the information sent to the processor. The information may be instructions for the processor, or data for the processor to use. A computer is not limited to just one processor; some machines require a vast volume of processing power and thus have multiple processors CPUs. In desktop computers, CPUs are mounted in a socket attached to the motherboard, and enclosed with a *heatsink* (as shown on the right) which helps to prevent overheating.

The processor speed is given in megahertz (MHz) or gigahertz (GHz). A single MHz means that the computer is performing one million instructions per second and a single GHz is one billion instructions per second. A processor can deal with a certain number of bits of information per instruction; with modern computers this is usually 64 bits, or 32 bits for a modern laptop.

Main memory

Main memory is the memory that is directly accessible to the processor. This is done in order not to slow the system down – an event known as *bottlenecking*. There are two types of main memory:

- *Read-only memory (ROM)* is a permanent area of storage for special programs installed during the process of computer manufacture. The contents are permanent because the data has been written onto a ROM chip which cannot be changed. Unlike RAM, ROM retains its contents even when the computer is turned off, being non-volatile, whereas RAM is volatile.

The *basic input/output system (BIOS)* is a part of ROM that stores critical information that starts the computer. For example, a pocket calculator contains 128 bytes of ROM to store the instructions of the calculator.

- *Random access memory (RAM)* constitutes the working area of the computer, used for storing programs and data currently in use. Modern RAM is measured in gigabytes (often gigahertz). Currently, the amount of RAM for a new modern desktop computer is around 4 to 6 gigabytes. Modern boards can support up to a staggering 64 gigabytes of RAM. RAM is *volatile*, meaning once the power supply is shut off all data that was stored is lost. RAM is directly written and read by the processor. This takes time, so the transfer rate the better; this is where frequency comes in.
- *Cache memory* holds frequently accessed code and data. It is extremely fast, close to the processor so that it doesn't need to access main memory where there is a bottleneck. This limits the size of how much data can be cached and so cache is just megabytes).

**COPYRIGHT
PROTECTED**



Did you know?!

Theoretically, the most RAM a 64-bit computer could possibly contain would be with today's technology and using modern standards of RAM capacity and main slow operations due to finding the specific RAM chip a piece of data was stored roughly 16 square kilometres!

Buses

The concept behind the term computer *bus* is similar to that of a vehicle bus – one place to another inside the computer. Although it is convenient to think of data being physically transferred from one place to another, in fact a bus line is just an (or many) wires along which 1s and 0s are sent. This word size is what limits the amount of data that can be transferred, called its *word size*.

As with most components, more is better – the larger a word size or bus is, the faster the computer. This is because the larger a computer has a larger bus size, it can transfer more data faster; it can reference larger numbers allowing more memory, so the computer can handle more and vary the number of instructions.

Buses are often common pathways shared by electronic signals to and from components in a computer. However, not all buses in the computer are the same and manufacturers use different buses; these are:

- *Data buses* carry the data being exchanged around a system in their entirety. They are likely to have the largest word size as actual information will be the data that needs to be sent via the bus. The other buses will determine how and where the data is to be processed or stored.
- *Address buses* carry the information about where the data is being sent. The same wires go to each component in turn. The components watch the address bus that they recognise is present. When this occurs the computer either reads data from the bus or places new data onto the data bus for the CPU to use. The number of unique addresses is equivalent to 2^n where n is the bus size.
- *Control bus* is the bus that carries the signal that regulates data flow. Operations such as memory writing, memory reading and bus arbitration follow strict timing sequences, some operations taking longer than others.
- *PCI (Peripheral Component Interconnect) bus* is the bus to which most components in a modern PC are connected. It runs at 33 MHz, with a bandwidth of 133 MBps, much faster than the old ISA standard, which only allows 8 Mbps. Finally, the PCI bus has a 32-bit width, compared to the ISA's 16-bit width.

I/O controllers

Input and output controllers provide the interface between the ports and the computer. Most modern motherboards have a number of ports built in so that the need for expansion cards, and controllers, are available for each of the port types below in order to add functionality to a computer.

INSPECTION COPY

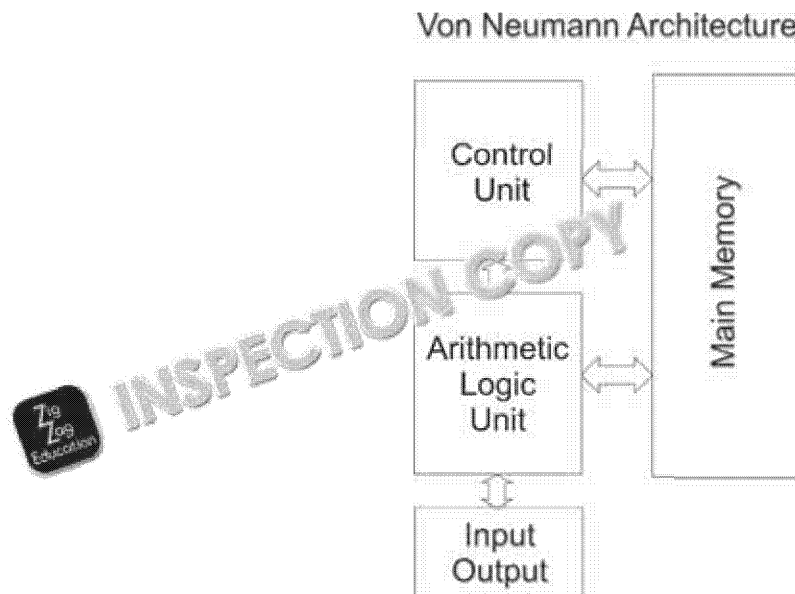
COPYRIGHT
PROTECTED



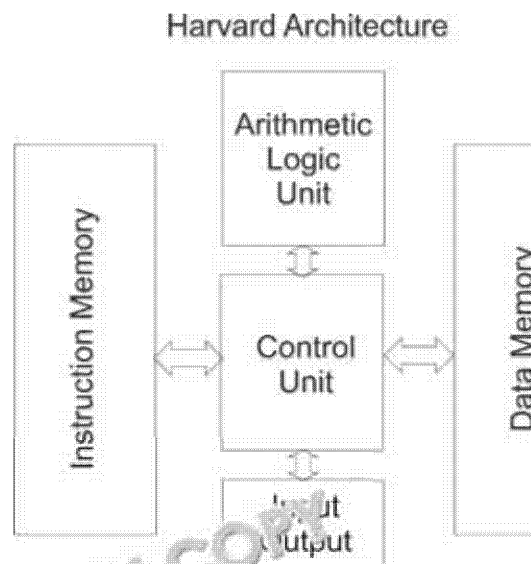
Computer architectures

There are two types of architecture you are expected to understand: the Von Neumann and the Harvard.

The Von Neumann was the earlier version in which memory is used to store both programs and data. The processor uses the fetch–decode–execute method (covered in 7.3) and therefore has a single path for both programs and data.



The Harvard architecture uses separate memory for data and programs/instructions. It also allows the processor to *pipeline*, often utilising a RISC processor.



A simple comparison summary is shown in the table below.

Von Neumann	Harvard
Single storage system for programs and data	Separate storage systems for programs and data
Each instruction takes two clock cycles (decode and execute)	Processor can complete an instruction in one clock cycle if pipelining is used
Pipelining cannot be implemented	Pipelining can be implemented
Older than Harvard, much more robust	Modern architecture

Embedded systems such as those found in digital signal processing systems use the Harvard architecture widely, whereas the Von Neumann architecture is primarily used in general-purpose computers.

INSPECTION COPY

**COPYRIGHT
PROTECTED**



Addressable memory

In the early days of computing, computers didn't need large volumes of RAM. Computers themselves were very simple and the majority of data processing could be done in the processor itself to complete the operation. However, as computers became more complex, they relied on the computer being able to store quantities of data on a medium and not just the currently being used.

Therefore, *addressable memory* is memory that is accessible from a computer (RAM). When the processor needs to access a section of memory it conveys the address over the address bus and uses the control bus to state whether it is reading from or writing to the memory. The data bus is then used to transmit data to or from the address location.

Questions: Internal Hardware Components

- 1 How are the buses used when transferring data? (3 marks)
- 2 What is RAM and what is it used for and why is it called 'volatile'? (3 marks)
- 3 What components are the following describing?
 - a) A component that contains the majority of the buses and acts as a central hub. All data must connect through. (1 mark)
 - b) A small volume of memory that contains frequently used data. This has a very limited volume. (1 mark)
 - c) A component that acts as the interface which all external data being transferred. Its transfer speed is measured in hertz which measures the number of calculations per second. (1 mark)

7.2 THE STORED PROGRAM CONCEPT

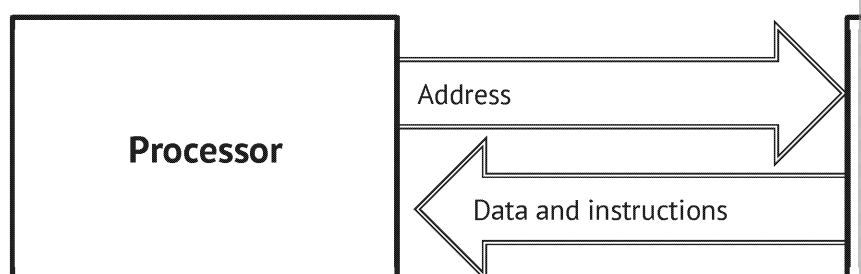
The stored program concept was born in the 1940s out of an idea of John Von Neumann. He proposed that computer programs could be run in computer memory rather than externally. This meant that the instructions to be run immediately and to be modified by the computer could be stored in memory.

The concept proposed that:

- In order to execute a program it must be contained in main memory.
- Instructions are read in machine code and are fetched, decoded if necessary, and then executed in the processor.
- The processor then performs the arithmetic, or logical, operations of the instructions.

The instructions would be collected in a periodic manner (serially) and then the processor would perform the arithmetic or logical operations that each instruction defined. The results of these operations would determine changes to the stored instructions and hence the course that program would take.

This concept allowed for the introduction of much more complicated and useful programs, paving the way for the type of computers that we are used to today.



**COPYRIGHT
PROTECTED**



This idea is linked closely with computer organisation and how we design a concept allowed digital computers to become more adaptable and more flexible concept. It was the first presentation of a computational device where the re could change the outcome of the program that required no human intervention.

The concept is still used to this day by most processor and computer manufacturers makes the machine. A computer can be adapted and built upon to perform any task to how the hardware will interact; without the stored program model there would be no what hardware could be used with other hardware which would make building computers which possess moderate computation capabilities, almost impossible.

Did you know?!

The introduction of the stored program concept as well as theoretical computer innovation in the field that has allowed us to develop our technology to where we are today and potential that it is, let's imagine what computer architectures would be like. The most notable of the stored program implementations is the Von Neumann architecture, covered in an earlier chapter.

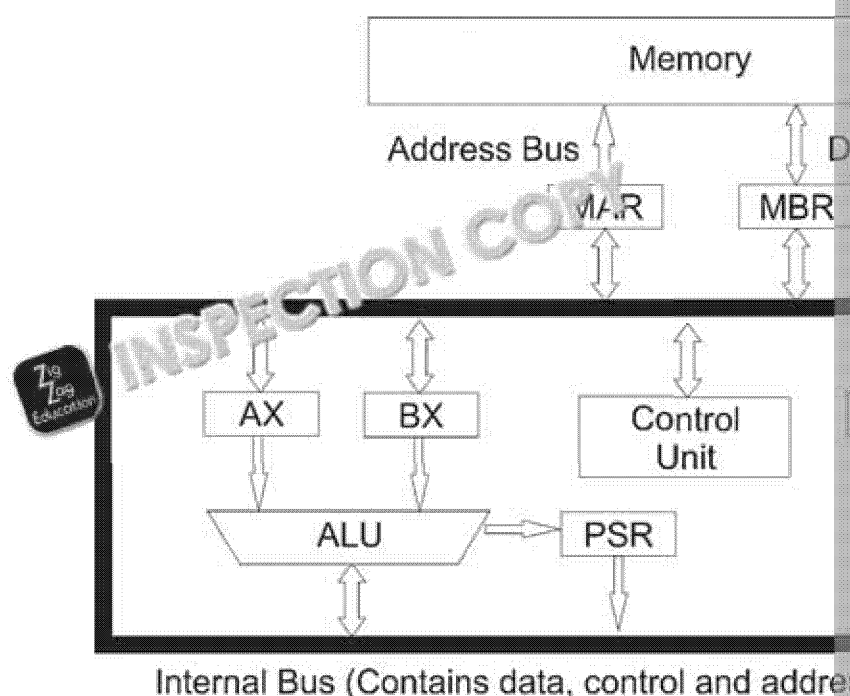
7.3 STRUCTURE AND ROLE OF THE PROCESSOR AND ITS COMPONENTS

THE PROCESSOR AND ITS COMPONENTS

The processor is a CPU (central processing unit) on a chip and provides the central processing. The specification of processors changes fast because there is a constant demand for more powerful processors.

The processor can be considered to be similar to a travel system in which data moves between different locations via buses. It is constantly changing the paths to the different elements and this is performed by the control unit and regulated by the system clock.

Inside a processor there are certain elements which are required in order to perform the instructions. The structure and role of the processor are covered in more detail on the following page.



COPYRIGHT
PROTECTED



Arithmetic logic unit (ALU)

Arithmetic logic unit carries out arithmetic operations such as addition, multiplication. It can also make logical comparisons between items of data; for example, it can check if one item is greater than another. Such logical operations can also be performed on non-numbers. The first Intel processor to have two ALUs so it could manipulate two sets of numbers. One of the ALU is held in a register called the accumulator.

Control unit (CU)

The *control unit* governs the operation of all hardware, including input and output devices. It does this by fetching, interpreting and executing each instruction in turn, in an automatic sequence. The fetch-execute cycle is described in detail below.

Clock

The *clock* is the part of the processor that regulates all of the actions that take place. It provides a regular pulse of high voltage followed by low voltage (voltage is the electric potential difference). Each high to low transition is known as a cycle, and each cycle implements a specific action.

Program counter (PC)

The *program counter*, sometimes called the *accumulator*, holds the address of the next instruction to be executed. When a sequence of instructions is to be executed the PC is automatically incremented by the address of the next instruction. Depending on the length of the current instruction, 1, 2 or 3 bytes are added to the current instruction is a jump, in which case the destination address is used.

General-purpose registers (AX, BX, CX and DX)

The other *general-purpose registers* are used for performing general arithmetic operations. They have no pre-defined role by the chip designer and as a result can be manipulated to perform the role of any other registers.

Memory buffer register (MBR)

Values about to be added or subtracted can be copied, via the *memory buffer register* (MBR), into the accumulators. The arithmetic result can be copied from there into a main memory location. All communications between the CPU and main memory take place through the MBR.

Memory address register (MAR) and current instruction register (CIR)

In order to fetch an instruction from memory the CPU places the address of the instruction into the *address register* and then carries out a memory read; the instruction is then copied into the *current instruction register*. Similarly, an instruction which itself requires data causes the address of the data word to be placed into the MAR. The execution of this instruction results in the copying of the address data word into the MBR, from where it can be accessed. The MBR acts as the point of transfer for both data and instructions passing between the main memory and the CPU.

Status or flag register (SR or FR)

The *status or flag register*, also known as the *processor status register* (PSR), contains flags based on the result of an instruction. These flags determine the operation of the next instruction. The result of the previous action. It also handles interrupts to signal to the controller.

**COPYRIGHT
PROTECTED**



THE FETCH-EXECUTE CYCLE AND THE ROLE OF REGISTERS

The control unit (CU) in the processor manages the execution of instructions in a sequence, decodes and synchronises it before executing by sending control signals to the computer. This is known as the *fetch-execute cycle* (or *fetch-decode-execute cycle*) in more detail below:

Fetch phase

Common to all instructions:

1. The contents of the PC are copied into the MAR. The MAR now contains the instruction and a memory read is initiated to copy the instruction word from memory.
2. The PC is incremented and now contains the address of the next instruction.
3. The instruction word is then copied from the MBR (MDR) into the CIR. The instruction is split into two parts: the operation code (*opcode*) and *operand*. The opcode is the instruction that tells the processor what operation to perform on the data to complete the operation on. *This is described further on the following page.*

Execute phase

The action taken is unique to the instruction:

1. The opcode instruction in the CIR is decoded to a simple operation such as *load* (see p.10) which affects the path the data will then follow.
2. The instruction in the CIR is executed; if the result needs to be compared or stored, it is held in the MAR.
3. Unless the instruction is a STOP instruction, then the cycle is repeated.

Example of the fetch-decode-execute cycle

Imagine a very simple computer that could have a program that consists of a few instructions. It might look like this:

```
LDA 2 (load 2 into the ALU)
ADD 1 (add 1 to 2 in the ALU - result in the accumulator)
STA result (store the result in memory)
HLT
```

In machine code these could be translated into instructions. Imagine a 2-bit instruction format where 11 = STA store in memory location and 00 = HLT; the last three bits hold the data.

```
01 010 LDA 2
10 001 ADD 1
11 101 STA (memory location 101 or 100)
00 000 HLT
```

This program is loaded into memory and stored in the locations as shown:

Location	Contents	Comment
001	01 010	Load 2 into ALU
010	10 001	Add 1 to the ALU result in the accumulator
011	11 101	Store accumulator in location 101 (memory location 101)
100	00 000	HLT - stop the processor
101	Empty	Where the result will be stored

INSPECTION COPY

COPYRIGHT
PROTECTED



The fetch–decode–execute cycle would do the following:

1. The PC is set to 001 which is sent to the MAR – this goes to the memory (010) into the MBR. The PC is incremented for the next instruction; the PC now has a value of 010.
2. The contents of the MBR (01 010) are transferred to the CIR; the CIR breaks down into the opcode (01) and the operand (010). The instruction is then executed and the result is stored in the ALU. Instruction complete.
3. The next instruction from the PC (010) is transferred to the MAR – this goes to the memory (10 001) into the MBR. The PC is incremented for the next instruction; the PC now has a value of 101.
4. The contents of the MBR (10 001) are transferred to the CIR; the CIR breaks down into the opcode (10) and the operand (001). The instruction is then executed and the result is stored in the accumulator which now has a value of (011) or 3, in decimal.
5. The next instruction from the PC (101) is transferred to the MAR – this goes to the memory (11 101) into the MBR. The PC is incremented for the next instruction; the PC now has a value of 110.
6. The contents of the MBR (11 101) are transferred to the CIR; the CIR breaks down into the opcode (11) and the operand (101). The instruction is then executed. In this case it is a store instruction and the contents of the accumulator are copied into that local memory location 101. Instruction complete.
7. The next instruction from the PC (110) is transferred to the MAR – this goes to the memory (00 000) into the MBR. The PC is incremented for the next instruction; the PC now has a value of 111.
8. The contents of the MBR (00 000) are transferred to the CIR; the CIR breaks down into the opcode (00) and the operand (00). The instruction is then executed. In this case it is HLT (halt) and the computer stops.

This is an example of a very small computer with a very limited instruction set. A larger instruction set as described in the next section.

THE PROCESSOR INSTRUCTION SET

These are the set of instructions that a processor can apply to a flow of data. All processors will have a different instruction set as there is not a standard instruction set. This gives manufacturers the capability to push their hardware further and faster. Modern instruction sets consist of roughly 100–250 instructions and cover a wide range of operations. If an instruction is read that the hardware does not expect, it will be required by the hardware, therefore if an instruction is read that is not in the set then an error flag is produced and the computer halts that program tree.

There are two approaches used in instruction sets:

- *Reduced Instruction Set Computer (RISC)*
- *Complex Instruction Set Computer (CISC)*

RISC vs CISC

RISC provides a fairly simple basic instruction set where in every operation the data is stored from memory to memory, whereas in CISC the instruction set is much more varied and can perform a wide range of operations across multiple clock cycles.

After the conception of the instruction sets, RISC tended to run faster because of its simplicity but it slowly became less commonplace because of how CISC machines were designed. The leader for CISC components is currently Intel and the reason why CISC became more common was because of how they managed to implement RISC principles into CISC architecture.

As the name implies, the instruction set is composed of instructions of a given type. It is divided into two parts: the *opcodes* and the *operand*. The *opcodes* (an abbreviation of *operation codes*) are the machine language instructions that tell the processor what operation is to be performed on the data that is stored in the operand address.

**COPYRIGHT
PROTECTED**



In a simple model, 4 bits might be assigned to the opcodes (3 bits for the basic operation to be done and 1 for the addressing mode) and 4 bits will be allocated to contain the address. The size of each instruction is 2 bytes. This means that a total of 16 instructions can be stored. For example, this might be the opcodes for addition using direct addressing.

ADDRESSING MODES

Direct addressing

Direct addressing specifies the actual or effective memory address containing the operand. This addressing mode allows for the code to be executed quickly and efficiently; it is also the most intuitive of the addressing modes to think about tasks, but it isn't always the best choice for a computer because it requires the processor to be processed and reallocated a memory address. The addressed memory location is used to obtain the operand.

An example of direct addressing is that if the command was STA 5, the 5 would be the address of the data. In the example of the simple computer in the fetch-decode-execute cycle, the instruction 11 101 (STA 5) is using direct memory addressing.

Immediate addressing

Immediate addressing is where the data is actually part of the command; for example, if the command was LDA 2, the processor will load the value 2 into the ALU register, not the contents of memory location 2.

MACHINE-CODE/ASSEMBLY-LANGUAGE OPERATIONS

The basic machine-code operators

The syllabus states that you should be able to understand and write programs in machine code, including immediate and direct addressing. There are various free versions of emulators such as Little Man Computer which will allow you to practise these skills as well as writing programs.

Three basic machine-code operations are LOAD, ADD and STORE. These are the operations that move data from memory to be moved about and added up.

LOAD and STORE are data transfer functions that make it possible to move data from memory and then store them in a relevant memory slot. ADD is an example of an arithmetic operation that takes a value from the memory and adds another number to it. Other arithmetic operations include SUBTRACT.

A typical machine-code segment might be LOAD A, [15] which would load the value from memory at location 15. Another would be ADD A, [16] which would add the value from memory at location 16 to the accumulator. So a simple program might consist of loading three numbers into the accumulator, adding them together and storing the result in the computer memory.

This program would look like as follows:

```
LOAD A, [15]
ADD A, [16]
ADD A, [17]
STORE A, [18]
Halt
```

In order for this to work, the instructions for this program (i.e. the code written in locations other than 15, 16, 17 and 18, and then the values from locations 15, 16 and 17) must be stored in memory. If the instructions were not saved in memory, the program would overwrite the instructions and hence destroy itself.

**COPYRIGHT
PROTECTED**



Operation	Example	Explanation
Load	Load A, [x]	Loads the contents of a variable, or a memory location into the accumulator so that it can be used in a calculation. X is usually found inside square brackets. X can be a memory location, the keyboard or device.
Add	Add A, [4]	Performs an addition operation on the contents of the accumulator and the value of a memory location, the keyboard or device.
Subtract	Subtr A, [4]	Similar to the Add command, performs a subtraction operation on the content of the accumulator. You can subtract the value of a memory address, or user input.
Store	Store A, [4B]	Stores the contents of the accumulator into a memory location. The memory location must be correctly defined.
Bitwise logic	OP A, B	The AND, OR, NOT and XOR commands perform bitwise logic operation A, B. These are used in the execution of them in other programming languages.
Compare	Comp A, [x >1]	Compares the contents of the accumulator with the value of a memory location to produce a True or False result. This is used to allow branching and determining the flow of a program.
Shift Left	SHL A,3	Multiplies a number by powers of 2 (2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768, 65536, 131072, 262144, 524288, 1048576, 2097152, 4194304, 8388608, 16777216, 33554432, 67108864, 134217728, 268435456, 536870912, 1073741824, 2147483648, 4294967296, 8589934592, 17179869184, 34359738368, 68719476736, 137438953472, 274877906944, 549755813888, 1099511627776, 2199023255552, 4398046511104, 8796093022208, 17592186044416, 35184372088832, 70368744177664, 140737488355328, 281474976710656, 562949953421312, 1125899906842624, 2251799813685248, 4503599627370496, 9007199254740992, 18014398509481984, 36028797018963968, 72057594037927936, 144115188075855872, 288230376151711744, 576460752303423488, 1152921504606846976, 2305843009213693952, 4611686018427387904, 9223372036854775808, 18446744073709551616, 36893488147419103232, 73786976294838206464, 147573952589676412928, 295147905179352825856, 590295810358705651712, 1180591620717411303424, 2361183241434822606848, 4722366482869645213696, 9444732965739290427392, 18889465931478580854784, 37778931862957161709568, 75557863725914323419136, 151115727451828646838272, 302231454903657293676544, 604462909807314587353088, 1208925819614629174706176, 2417851639229258349412352, 4835703278458516698824704, 9671406556917033397649408, 19342813113834066795298816, 38685626227668133590597632, 77371252455336267181195264, 154742504910672534362390528, 309485009821345068724781056, 618970019642690137449562112, 1237940039285380274899124224, 2475880078570760549798248448, 4951760157141521099596496896, 9903520314283042199192993792, 19807040628566084398385987584, 39614081257132168796771975168, 79228162514264337593543950336, 158456325028528675187087900672, 316912650057057350374175801344, 633825300114114700748351602688, 1267650600228229401496703205376, 2535301200456458802993406410752, 5070602400912917605986812821504, 10141204801825835211973625643008, 20282409603651670423947251286016, 40564819207303340847894502572032, 81129638414606681695789005144064, 162259276829213363391578010288128, 324518553658426726783156020576256, 649037107316853453566312041152512, 1298074214633706907132624082305024, 2596148429267413814265248164610048, 5192296858534827628530496329220096, 10384593717069655257060992658440192, 20769187434139310514121985316880384, 41538374868278621028243970633760768, 83076749736557242056487941267521536, 166153499473114484112975882535043072, 332306998946228968225951765070086144, 664613997892457936451903530140172288, 1329227995784915872903807060280344576, 2658455991569831745807614120560689152, 5316911983139663491615228241121378304, 10633823966279326983230456482242756608, 21267647932558653966460912964485513216, 42535295865117307932921825928971026432, 85070591730234615865843651857942052864, 170141183460469231731687303715884105728, 340282366920938463463374607431768211456, 680564733841876926926749214863536422912, 1361129467683753853853498429727072845824, 2722258935367507707706996859454145691648, 5444517870735015415413993718908291383296, 10889035741470030830827987437816582766592, 21778071482940061661655974875633165533184, 43556142965880123323311949751266331066368, 87112285931760246646623899502532662132736, 174224571863520493293247799005065324265472, 348449143727040986586495598010130648530944, 696898287454081973172991196020261297061888, 1393796574908163946345982392040522594123776, 2787593149816327892691964784081045188247552, 5575186299632655785383929568162090376495104, 11150372599265311570767859136324180752990208, 22300745198530623141535718272648361505980416, 44601490397061246283071436545296723011960832, 89202980794122492566142873090593446023921664, 178405961588244985132285746181186892047843328, 356811923176489970264571492362373784095686656, 713623846352979940529142984724747568191373312, 1427247692705959881058285969449495136382746624, 2854495385411919762116571938898990272765493248, 5708990770823839524233143877797980545530986496, 11417981541647679048466287755595961091061972992, 22835963083295358096932575511191922182123945984, 45671926166590716193865151022383844364247891968, 91343852333181432387730302044767688728495783936, 182687704666362864775460604089535377456991567872, 365375409332725729550921208179070754913983135744, 730750818665451459101842416358141509827966271488, 1461501637330902918203684832716283019655932542976, 2923003274661805836407369665432566039311865085952, 5846006549323611672814739330865132078623730171904, 11692013098647223345629478661730264157247460343808, 23384026197294446691258957323460528314494920687616, 46768052394588893382517914646921056628989841375232, 93536104789177786765035829293842113257979682750464, 187072209578355573530071658587684226515959365500928, 374144419156711147060143317175368453031918731001856, 748288838313422294120286634350736906063837462003712, 1496577676626844588240573268701473812127674924007424, 2993155353253689176481146537402947624255349848014848, 5986310706507378352962293074805895248510699696029696, 11972621413014756705924586149611790497021399392059392, 23945242826029513411849172299223580994042798784118784, 47890485652059026823698344598447161988085597568237568, 95780971304118053647396689196894323976171195136475136, 191561942608236107294793378393788647952342390272950272, 383123885216472214589586756787577295904684780545900544, 766247770432944429179173513575154591809369561091801088, 1532495540865888858358347027150309183618739122183602176, 3064991081731777716716694054300618367237478244367204352, 6129982163463555433433388108601236734474956488734408704, 12259964326927110866866776217202473468949912977468817408, 24519928653854221733733552434404946937899825954937634816, 49039857307708443467467104868809893875799651909875269632, 98079714615416886934934209737619787751599303819750539264, 196159429230833773869868419475239575503198607639501078528, 392318858461667547739736838950479151006397215279002157056, 784637716923335095479473677900958302012794430558004314112, 1569275433846670190958947355801916604025588861116008628224, 3138550867693340381917894711603833208051177722232017256448, 6277101735386680763835789423207666416102355444464034512896, 12554203470773361527671578846415332832204710888928069025792, 25108406941546723055343157692830665664409421777856138051584, 50216813883093446110686315385661331328818843555712276103168, 100433627766186892221372630771322662657637687111424552206336, 200867255532373784442745261542645325315275374222849104412672, 401734511064747568885490523085290650630550748445698208825344, 803469022129495137770981046170581301261101496891396417650688, 1606938044258990275541962092341162602522202993782792835301376, 3213876088517980551083924184682325205044405987565585670602752, 6427752177035961102167848369364650410088811975131171341205504, 12855504354071922204335696738729300820177623950262342682411008, 25711008708143844408671393477458601640355247900524685364822016, 51422017416287688817342786954917203280710495801049370729644032, 102844034832575377634685573909834406561420991602098741459288064, 205688069665150755269371147819668813122841983204197482918576128, 411376139330301510538742295639337626245683966408394965837152256, 822752278660603021077484591278675252491367932816789931674304512, 1645504557321206042154969182557350504982735865633579863348609024, 3291009114642412084309938365114701009965471731267159726697218048, 6582018229284824168619876730229402019930943462534319453394436096, 13164036458569648337239753460458804039861886925068638906788872192, 26328072917139296674479506920917608079723773850137277813577744384, 52656145834278593348959013841835216159447547700274555627155488768, 105312291668557186697918027683670432318895095400549111254310977536, 210624583337114373395836055367340864637790190801098222508621955072, 421249166674228746791672110734681729275580381602196445017243910144, 842498333348457493583344221469363458551160763204392890034487820288, 1684996666696914987166688442938726917102321526408785780068975640576, 3369993333393829974333376885877453834204643052817571560137951281152, 6739986666787659948666753771754907668409286105635143120275902562304, 13479973333575319897333507543509815336818572211270286240551805124608, 26959946667150639794667015087019630673637144422540572481103610249216, 53919893334301279589334030174039261347274288845081144962207220498432, 107839786668602559178668060348078522694548577690162289924414440996864, 215679573337205118357336120696157045389097155380324579848828881993728, 431359146674410236714672241392314090778194310760649159697657763987456, 862718293348820473429344482784628181556388621521298319395315527974912, 1725436586697640946858688965569256363112777243042596638790631055949824, 3450873173395281893717377931138512726225554486085193277581262111899648, 6901746346790563787434755862277025452451108972170386555162524223799296, 13803492693581127574869511724554050904902217944340773110325048447598592, 27606985387162255149739023449108101809804435888681546220650096895197184, 55213970774324510299478046898216203619608871777363092441300193790394368, 110427941548649020598956093796432407239217743554726184882600387580788736, 220855883097298041197912187592864814478435487109452369765200775161577472, 441711766194596082395824375185729628956870974218904739530401550323154944, 883423532389192164791648750371459257913741948437809479060803100646309888, 1766847064778384329583297500742918515827483896875618958121606201292619776, 3533694129556768659166595001485837031654967793751237916243212402585239552, 7067388259113537318333190002971674063309935587502475832486424805170479104, 14134776518227074636666380005943348126619871175004951664972849610340958208, 28269553036454149273332760011886696253239742350009903329945699220681916416, 56539106072908298546665520023773392506479484700019806659891398441363832832, 113078212145816597093331040047546785012958969400039613319782796882727665664, 226156424291633194186662080095093570025917938800079226639565593765455331328, 452312848583266388373324160190187140051835877600158453279131187530910662656, 904625697166532776746648320380374280103671755200316906558262375061821325312, 1809251394333065553493296640760748560207343510400633813116524750123642650624, 3618502788666131106986593281521497120414687020801267626233049500247285301248, 7237005577332262213973186563042994240829374041602535252466099000494570602496, 14474011154664524427946373126085988481658748083205070504932198000989141204992, 28948022309329048855892746252171976963317496166410141009864396001978282409984, 57896044618658097711785492504343953926634992332820282019728792003956564819968, 115792089237316195423570985008687907853269984665640564039457584007913129639936, 231584178474632390847141970017375815706539969331281128078915168015826259279872, 463168356949264781694283940034751631413079938662562256157830336031652518559744, 926336713898529563388567880069503262826159877325124512315660672063305037119488, 1852673427797059126777135760139006525652319754650249024631321344126610074238976, 3705346855594118253554271520278013051304639509300498049262642688253220148477952, 7410693711188236507108543040556026102609279018600996098525285376506440296955904, 14821387422376473014217086081112052205218558037201992197050570753012880593911808, 29642774844752946028434172162224104410437116074403984394101141506025761187823616, 59285549689505892056868344324448208820874232148807968788202283012051522375647232, 118571099379011784113736688648896417641748464297615937576404566024103044751294464, 237142198758023568227473377297792835283496928595231875152809132048206089502588928, 474284397516047136454946754595585670566993857190463750305618264096412179005177856, 948568795032094272909893509191171341133987714380927500611236528192824358010355712, 1897137590064188545819787018382342682267975428761855001222473056385648716020711424, 3794275180128377091639574036764685364535950857523710002444946112771297432041422848, 7588550360256754183279148073529370729071901715047420004889892225542594864082845696, 15177100720513508366558296147058741458143803430094840009779784

ARM assembly language

The AQA specification states that the code written in exams will be of the form where the commands will be described in the paper.

ARM technology uses RISC technology and is used in devices such as the Raspberry Pi and mobile phones. The ARM registers and instruction sets are different from those for x86 Intel processors as memory and instructions are stored and referenced as separate locations.

The basic differences are described below and include additional command sets. ARM technology also has more registers to be used within the processor, and it is recommended that these are used rather than linking to memory as this simplifies the process. Although registers can be copied to memory (using direct or indirect memory addressing), this can significantly slow down the program. For most examples, utilising the registers r4 to r9 (as variables / general purpose registers) is preferred. Although commands in the Little Man Computer and ARM may look similar, their syntax is different. For example, in the Little Man Computer instruction for ADD is `ADD num1` with the result stored in the accumulator and store the result in the accumulator. In ARM, the instruction for ADD is `ADD r0, r4, r5` which would add the contents of r5 to the contents of r4 and store the result in r0.

Example showing the difference between Little Man Computer and ARM

If we were to write a program to take in two numbers and output the result, in the Little Man Computer it would be done using the routines `inp` and `out` inside an emulator whereas the Raspberry Pi (RPi) uses its own registers, so the program is more complex due to setting up the messages and output but the fundamental assembly is the same. The Little Man Computer also utilises the routines `bl` and `printf` to output which allows the output to be formatted.

LMC code to add two numbers and output the result

```
INP
STA num1
INP
STA num2
ADD num1
OUT
```

ARM code to add two numbers and output the result for a RPi ARM

(note: included statements for input)

```
@ set up data reference and output statements
.data
.align 2
scan_format:
.asciz "%d"
.align 2
out_format:
.asciz "The sum is : %d\n"
.align 2
instr1:
.asciz "Enter first integer."
instr2:
.asciz "Enter second integer."
.align 2
num1:
.word 0
num2:
.word 0
.text
.global main
main:
```

INSPECTION COPY

COPYRIGHT
PROTECTED



```

push {ip, lr}           @ used with pop at end of main, allow
                        @ of the program
ldr r0, =instr1         @ load the instruction into r0
bl puts                 @ output to screen
ldr r1, =num1           @ sets up num1 for input
ldr r0, =scan_format    @ calls routine to input num1
bl scanf               @ load the instruction into r0
ldr r0, =instr2
bl puts
ldr r1, =num2           @ sets up num2 for input
ldr r0, =scan_format    @ calls routine to input num2
bl scanf
ldr r6, =num1           @ load address of num1 into r6 needed
ldr r4, [r6]            @ load value of num1 into r4 nb other
                        @ allow ldr r4,num1

ldr r6, =num2
ldr r5, [r6]
add r1, r4, r5          @ add contents of r5 to r4 and store
ldr r0, =out_format     @ ~ output for "answer is: " follow
bl printf               @ output the solution
pop {ip, lr}           @ used with push at start of main, al

```

Commands similar to Little Man Computer

Below is a table of the ARM type commands which are similar to Little Man C

Mnemonic	Example	Function
ADD	ADD r0,r4,r5	Addition adds contents to r5 to r4 and store
SUB	SUB r0,r4,r5	Subtraction subtracts r5 from r4 and stores i
RSB	RSB r4,r4,#300	Reverse subtraction subtracts r4 from secon 300) and stores it in this case back to r4
MUL	MUL r0,r4,r5	Multiplication multiplies r4 by r5 and stores answer in a register used in the calculation r4,r4,r5 will cause an error
AND	AND r0,r4,r5	Bitwise AND between r4 and r5 returned in
ORR	ORR r0,r4,r5	Bitwise OR between r4 and r5 returned in rC
EOR	EOR r0,r4,r5	Bitwise exclusive OR between r4 and r5 retu
MVN	MVN r0, r4	Bitwise NOT stored in r0
TST	TST r4,r5	Test (performs bitwise AND sets flags accord
CMP	CMP r4,r5	Compare r4 and r5 (actually uses SUB) but e values and sets flags to show whether they are e
B	B label	Unconditional jump/branch to a label
BEQ	BEQ label	Conditional jump/branch to a label based on
MOV	MOV r0, r4	Move contents of r4 to r0
LDR STR	ldr r3, =sum str r0, [r3]	Load and store work together, so in this exa The value in r0 is then stored in =sum using
BL	BL printf	Call a subroutine then return; in this exampl routine

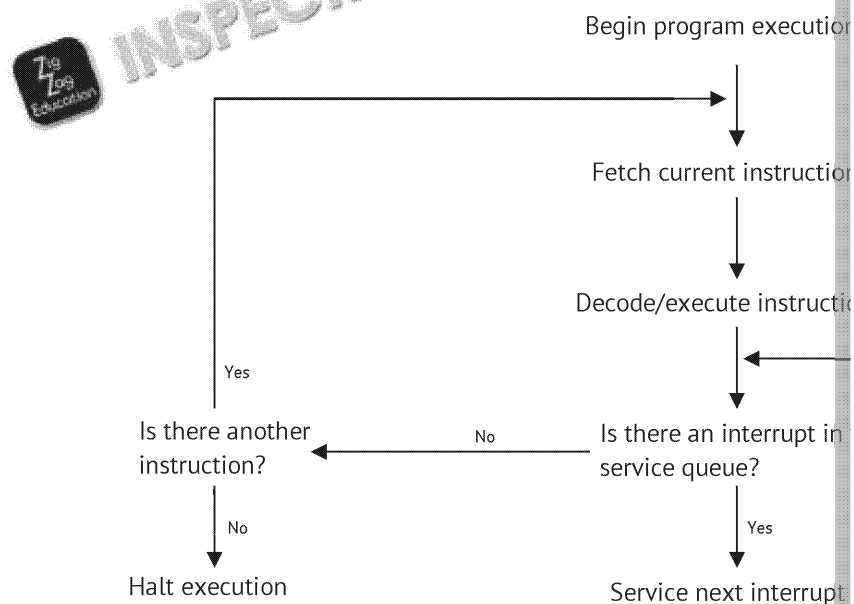
**COPYRIGHT
PROTECTED**



INTERRUPTS

Interrupts are the computer's way of diverting away from the sequential nature of program execution. Without interrupts a computer would continue to finish all the instructions for a program before checking to see whether there are any interrupts. When an internal error occurs within the computer, the processor is sent from the device to notify the processor. Interrupts are stored in a *priority queue*. In modern-day general-purpose computers the interrupt queue is checked after each instruction is executed. If the queue is empty then the next instruction in the sequence is executed. If there is something in the queue then the processor runs a program called the *interrupt service routine*.

Each type of interrupt has its own interrupt service routine, so the ISR for an internal error is different from the ISR for a paper jam in a printer. However, while the ISR is at work solving the problem, the volatile environment of memory was running before the interrupt occurred is stored so that when control is passed back to the main program the user can continue to use the computer.



Types of interrupt

As mentioned earlier, not all interrupts are the same; each interrupt needs its own service routine. To understand what needs to happen when that interrupt is present in the interrupt queue, you should be familiar with just by general use of a computer.

I/O interrupts

Input/output interrupts occur when the computer is performing a data transfer and the transfer has been completed or there has been an error during transmission. An example of this is the transfer of a file from main memory to a USB memory stick and during the transfer you are notified without warning. The completion generates an interrupt that prevents the further execution of the program as the destination address is no longer reachable.

Timer interrupts

There is a concept in computing and technology called *time-critical* data. This is data that is critical files that take priority over the current instruction being performed. For example, if you stream a movie over a service such as Netflix – if the data was not received in time, the audio and the video would be out of sync because the video portion of the data would be transmitted and processed faster than the audio. Another example of this can be seen in loading a YouTube video; the audio would play first while the video would be blank or only the first frame would be displayed until the video buffer caught up.

INSPECTION COPY

COPYRIGHT
PROTECTED

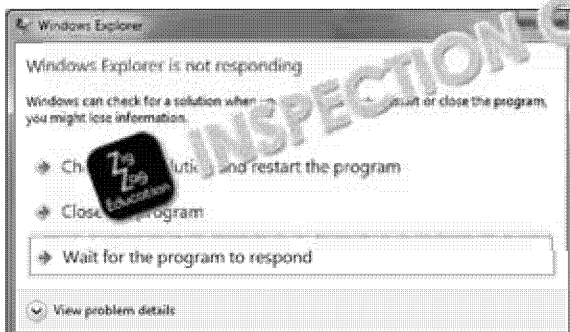


Hardware interrupts

These are the most common type of interrupt that all computer users will have experienced. These interrupts mark critical errors in the hardware of the computer itself that affect the operation of the OS. They are performed by the CPU. For example, if there is a sudden loss of power, the CPU will save as much data as possible and close the OS down as safely as possible to prevent data loss. The OS is then rebooted.

Program Interrupts

These interrupts are produced by the software of a computer system to tell the user that an event has occurred or an exception has been met. These frequently occur in games but will be handled by the system to not disrupt the gameplay experience.



The most common cause of a program hang produced is memory access. A program may access a memory location that is not available. When a program enters a state where it is unresponsive, which is what a hang is, the Windows operating system will attempt to terminate the program.

FACTORS AFFECTING PROCESSOR PERFORMANCE

Multiple cores

A processor that supports multiple *cores* is essentially a processor with multiple identical processors integrated into a single chip. The increase in cores allows for a greater throughput of data. It allows the processor to divide the labour of performing a task between all the cores.

If the software is written for multiple cores, the processor can split the instructions into a number of simpler instructions in a process called 'threading'. If the software is not coded for use of multiple cores then the extra cores aren't used for that process and are assigned for performing other background tasks.

Cache memory

The *cache* memory for a processor is a small amount of *high-performance* RAM. It enables the CPU to access repeatedly used data directly from its own board, requesting it from system memory. Cache is critical to applications such as video applications but is not critical for general applications such as emails and word processing.

Words

A *word* is a group of bits that can be addressed, transferred and manipulated by a processor. The size of a word, the word length, is determined by the width of the computer and is usually a multiple of a byte, thus consisting of 16, 24, 32 or 64 bits. A larger word size may mean that a computer can work faster than a computer with a smaller word size, but it is also more efficient when that data does not need to use the full word length.

Clock speed

Every computer contains an internal clock; this clock is used to regulate the instructions that are carried out. The central processing unit in the computer needs to have a clock to carry out an instruction; this means that the faster a computer's clock 'ticks', the more instructions are carried out per second. The speed of the clock is measured in megahertz (MHz).

Bus width

The width of the bus is another factor that relates to the performance of the computer. It is the amount of data that the central processing unit can transmit at a time to the memory or the input and output devices. For example, a 16-bit bus can carry 16 bits of data at a time. Consider the bus in terms of passengers; if each passenger was carrying some data to a destination you would want to fit as many on the bus as possible to save doing multiple trips.

Did you know?!

The ultimate reason for CPU performance comes down to how temperature affects the CPU. You may know that it has a large metal hood over the chip – this is actually a heat sink designed to carry heat away from the delicate components. If the temperature of the motherboard will reduce its clock speed, drop threads and reduce the number of instructions it can carry out. This is called thermal throttling and is the computer's way of cooling the CPU down. If the temperature gets too high, the CPU will simply shut itself off.

Overclocking is the process of increasing the clock speed of the processor through software. The faster the instructions can perform. However, as above, the temperature can become critical and so 'overclockers' may spend more money on different methods of keeping the CPU cool, such as water cooling, to prevent throttling.

Questions: Structure and the Role of the Processor

- 1 Which register in the processor must all input data pass through? (1 mark)
- 2 Describe all the steps required in the fetch and execute steps of the fetch-execute cycle. (4 marks)
- 3 Your computer is sat idle at the desktop; you are connected to the Internet but no programs are running.
 - a) Why do your CPU and memory usage never reach 0%? (2 marks)
 - b) You decide to load an application; are interrupts used? (1 mark)
- 4 Are each of the following statements true or false?
 - a) The entirety of a frequently used program is stored to the CPU's cache. (1 mark)
 - b) The word size is controlled by the CPU. (1 mark)
 - c) All software will automatically use multiple cores. (1 mark)
 - d) Increasing the bus width will improve performance. (1 mark)

**COPYRIGHT
PROTECTED**



7.4 EXTERNAL HARDWARE DEVICES

INPUT AND OUTPUT DEVICES

Barcode reader


Barcodes are used everywhere nowadays. They are generally used to track items, such as products in a supermarket, or books in a library. Although there are many different specific systems for encoding the information in a barcode, they all work in a similar manner. In the EAN and UPC barcode systems, the most popular for food in the UK and USA, two bars and two spaces represent each character in the code. The actual system used is quite complex because the scanners have to cope with the barcode being folded, or odd or varying speed. The two types of reader commonly found connected to a computer are the wand reader that reads across the barcode (wands) and the imager that reads the entire barcode at once.

How does a barcode reader work?

A barcode is made up of a series of bars which in combinations form the different characters. The bars are of a set distance and so some multiple bars appear as thicker lines. A laser light is directed across the barcode. In most cases this is a laser light which is directed at the barcode. The light will either be absorbed by the black bars or reflected by the white spaces. A photo diode (a device that produces an electrical signal based on the amount of light it receives) produces an electrical signal based on the amount of light it receives. This signal is interpreted as a 1 or 0 and fed into the computer. The combinations of 1s and 0s form the bar code number. The number is then used as an identifier on a computer database to obtain details about the item that has been scanned.

Radio frequency identification (RFID)

Radio frequency identification is used in a similar way to barcodes in that the details are stored in a computer database to obtain the details. However, the technology uses radio waves instead of a barcode.



that sends out an identifying signal that can be interpreted by a computer. RFID is used in many applications, including passports, animal identifications and libraries among many others. RFID can be passive or active; passive and semi-passive rely on the power to send a response, whereas active devices have their own power source and broadcast a signal continuously. Semi-passive devices can facilitate functions with the device (such as data storage).

Digital camera

Digital cameras and digital video cameras take pictures and store them digitally rather than on analogue or photographic film. The storage device is unimportant from a photography point of view, but has space (MP) implications and size (portable) implications. Most cameras these days use flash memory cards to store images on, with the top-of-the-range cameras having many gigabytes of memory, meaning they can store thousands of high-resolution images.

In addition, since the images are stored digitally, the camera can manipulate the images directly, applying effects such as sepia tone to the image, and easily and quickly editing and saving the image in a word-processing document. Of course, the real advantage is that you do not have to wait for the film to be developed after you take the picture.

How do digital cameras work?

The lens on a digital camera magnifies and focuses the light in the same way as a conventional camera, but instead of focusing the light onto a conventional film (which then creates a chemical image), digital cameras focus the light onto a device known as a charge-coupled device (CCD). This is an array of transistors which create electrical currents in proportion to the intensity of the light used for each pixel and for each primary colour. The image is then stored on a memory card.

INSPECTION COPY

COPYRIGHT
PROTECTED



Laser printer

These give a very high-quality print. For an example of running costs, the HP of paper from one new cartridge (which currently costs about £30) – a cost of Colour laser printers are relatively quick and high quality but can be expensive.

How do laser printers work?

1. The computer's printer driver sends information to the printer about
2. The printer uses a special wire called a 'corona wire' which statically even across the whole area.
3. The printer's processor uses the received information to shine a laser the printer drum. The laser does not actually print; the beam is directed. The drum becomes negatively charged in those places where the laser positively charged elsewhere.
4. The toner (a carbon powder) is positively charged with another corona
5. The drum rotates and the positively charged toner jumps onto the place charged. Equally, the toner is repelled by the positive charge, making the
6. Meanwhile, the paper is rolled through the paper train and given a strong The paper is then fed very close to the drum and the toner particles jump
7. Finally, the paper with its particles of toner are passed through two paper. This is known as fusing.

SECONDARY STORAGE DEVICES

What is secondary storage?

The most common peripheral is secondary storage, which in this context means stores information permanently and is not immediately accessible by the processor. The hard disk as being the main memory of the computer, in fact the computer takes programs stored on the hard disk than those stored in RAM. The computer copies from the hard disk to RAM, then writes back to the hard disk where required. (For a very large file or lots of programs running at once) it will start using the hard disk and the computer appears to slow down.

Hard disk drives (HDDs)

A HDD is a storage medium that stores large amounts of data. They first began in the 1970s as a method of quickly storing data; they were faster than the tape system and less expensive. For this reason, they were typically used to store program information, much like how RAM is used nowadays. Over the years, however, their cost dropped and they became a standard piece of computer equipment.

Nowadays, they are still used for storing programs, documents and anything else that needs to be stored. External hard drives are becoming increasingly popular as the amount of information that a typical computer uses is growing rapidly. External hard drives connect through USB and provide additional storage when the internal hard disks become full.

Information is held in blocks formed by tracks and sectors (see diagram). Each block of information contains the same amount; this means that information is more compact closer to the centre of the disk.

**COPYRIGHT
PROTECTED**



Rotation speed

A hard disk consists of one or more disks which spin very rapidly in an evacuated enclosure such that the reading heads, which are very small indeed, are suspended above the disks by a thin layer of air. This is known as the *levitation effect*, typically about two millionths of a centimetre for a modern hard drive. This is to prevent specks of dust, or other airborne detritus, from damaging the disk surface. The speeds of rotation (typically 7200 rpm).

Capacity

The early hard drives had a capacity of about 5 MB, but nowadays it is common to find specifications of above 1 TB of hard drive space. At the same time, the cost has fallen so much that hence hard drives are more affordable than ever. By way of example, a 10 MB hard drive might have cost £1,500 – a cost per megabyte of about £150. By contrast, in 2010, a 1 TB drive was purchased for as little as £50, giving a cost per megabyte of about 0.005p per MB.

Optical disk

CDs and DVDs are metal discs embedded in a plastic protective housing. Each disc is created by the process of creating the disk, and placing the data on it.

Although you can get re-writable (RW) versions, most optical disks in use are *(Write Once, Read Many)* media; this refers to the fact that once they have been written, there is no way to change the data on them. For this reason, and the fact that they are a de facto standard for the distribution of software and other media such as music and video, they are very popular.


More recently, Blu-Ray discs are becoming more widespread, mainly because of their large storage capacity (up to 50 GB) makes them a popular choice for storing data – such as content such as HD films and video games.

Optical spiral

Data on the disk is arranged in a very different way from that on a hard drive. On a hard drive, data is arranged in concentric tracks and sectors. For example on a CD-ROM, data is written in a single track from the inside of the CD to the outside. The packing density is uniform throughout the track in order to sustain transfer rates at the inside of the disc, with reference to the outside. Another difference between optical disks and hard drives is in the way they store data. Hard drives use a magnetic system, whereas CD-ROMs use a purely optical system.

Binary pits

Originally, during the mastering process, a very powerful laser was used to burn pits into the disc. When the disc is read back, the reading laser sees these pits as one sort of part of the track as the other. In this way, we have the storage of 0s and 1s for digital data storage. When the user reads the disk, a high intensity laser is shone at the surface. The pits and peaks produce a difference in how they scatter the light; this is picked up by a photodiode and converted into the binary digits. Music CDs and CD-R (for distribution of software) may still use this method as this makes the disk read-only.



**COPYRIGHT
PROTECTED**



Modern CD-RW/DVD-RW use a rewritable technology in the form of dye which is written by a laser beam. The dye has a special property whereby if it is heated to a certain temperature it becomes transparent, whereas if it is cooled to a lower temperature it becomes opaque. This feature allows the disk to be reused. The principle of the rewritable technology is the same as the principle of the read-only technology, with the transparent layer acting like a pit.

Solid-state drive (SSD)

Flash drives are everywhere in modern technology; something that was once considered a luxury is now being reduced to the point where two or more modules can fit inside a hard drive. Flash drives are a type of memory known as Not-And (NAND) flash memory that uses a different method of reading and writing of data. Modern operating systems find it easier to transfer data in blocks as *pages*; a page is a fixed block of memory that is used to facilitate the efficient transfer of data. The controller checks for certain constraints. Firstly, when you combine many pages to form a block – but a block cannot overwrite other information on the pages. Secondly, before you can write to an allocated memory location you must first erase the page that has been allocated to that location, but the underlying technology requires that the entire block is erased if the page is linked to a block. The result is a form of memory that has very low latency (response time) and transfer speeds that vastly outpace that of the magnetic or optical format.

These benefits do come at a price, however. The way in which the data is handled is damaging for the memory components of the SSD and it can lead to data errors. SSDs have an expected lifetime that is based on the number of read/write cycles that they can handle; the more the data on the drive is accessed, the more the lifetime is decreased. Recent studies have shown that SSDs last longer than expected but the cost still remains high compared to magnetic hard drives. That most who employ a solid-state drive only store the OS and frequently accessed files, while storing all other personal files to a magnetic drive and keeping the SSD as a cache.

Speeds of access and suitability

The suitability of which medium you use depends entirely on your requirements. Technology is changing in the way that data is stored even in the last 10 years, and there are new technologies being made every day so what is right now may change in the near future. It is easy to see why hard drives (HDDs) have been employed all over the world in modern computers. They have a large storage capacity and have a relatively fast read and write speed for personal use. Their primary role is to facilitate the storage of your personal files, the OS to interact with the user.

Solid-state drives, while having fast transfer speeds, are expensive to produce. They are just beginning to match the capacity of HDDs but with a much higher price tag. LSI announced it had created a 4 terabyte solid-state drive that would retail for the same time a 4 terabyte HDD would cost, around £100.

Optical disks are slowly becoming the way of the floppy disk; they're becoming obsolete. Manufacturers are starting to produce computers that don't have optical bays. Access to data is slowly becoming digital; everything can just as easily be placed on a hard drive. There is a hierarchy of speeds for I/O devices; in descending order they are:

Device	Read/write speed
Solid-state	550 MB/s
Hard disk	300 MB/s
Optical disk	7,200 KB/s

**COPYRIGHT
PROTECTED**



8. Consequences of Uses of Computers

The power of modern computers and the volume of data handling has led to a whole new set of considerations. This area will continue to grow as technology continues to expand and increase. This section discusses moral consequences of computing, and the professional computer scientist must always be aware of the responsibilities that computing gives with responsibility.

This section covers:

Individual consequences.....	p1	Legal consequences.....	p1
Social consequences.....	p3	Cultural issues.....	p3

8.1 INDIVIDUAL, SOCIAL, LEGAL AND CULTURAL ISSUES

In the mid 1940s the smallest computer occupied 1,800 square feet, used over 10,000 vacuum tubes (a vacuum tube is a version of a transistor) and weighed nearly 50 tonnes. It was capable of performing 1,000 operations per second and consumed an amazing 150,000 watts of power. Modern computers, with their integrated circuits, are now able to fit neatly under a desk and even into the palm of your hand. The power of computers can now be applied to anything with an embedded microchip in it.

As technological breakthroughs are made it becomes more common place to find computers in every appliance found in a home or workplace. With great technological power comes great responsibility. Many people take into account when using computers. These consequences can be categorized into four main areas:

- *Individual* consequences cover the *moral* standings of computer use. This includes the physiological and physical effects of computer use as well as how you use the computer.
- *Social* consequences cover the *ethical* standings of computer use. This includes the impact on individuals, governing their principles and behaviour.
- *Legal* consequences are the *written* laws that govern computer use and the challenges facing legislators.
- *Cultural issues* can cover a large variety of meanings from traditions and beliefs to the impact of technology on society.

Note: you are not required to know the specific acts and laws governing social and legal consequences. In context it would be a good idea to look into these acts to build on your understanding. Understanding these matters will better your awareness of how you yourself use computers.

INDIVIDUAL CONSEQUENCES

Morality

Morality is the notion of what is right and wrong or good and evil in society. As technology has become larger and more widespread, there have been issues that we have needed to be aware of as individuals when using their computer. As most computer users are not as technically minded as the programmers, the majority of the power lies with computer scientists. However, those who develop the software also have the responsibility to ensure that the software is not malicious.

Did you know?!

When you download social media apps on to a smartphone there is a small term and conditions agreement you agree to before the app will download. Many social media state the requirements for using the app. When you agree to these you often unknowingly sign in agreement to the app having access to your messages, call history and, most worryingly, your camera and microphone. In fact, Cyanogen mod now comes with a built-in data access protector to prevent apps from accessing your data and remain private.

**COPYRIGHT
PROTECTED**



Health and Safety

Prolonged use of computers is a relatively new occurrence that has been studied. It has uncovered several health risks. Stress is common in employees where communication is needed due to the fast-paced nature of communication. There is a wide variety of disorders that are linked to computer use. These pressures can make you more susceptible to health issues, potentially making it dangerous in the long term.

When working with computers, workers are at risk from health issues that may arise. One of the most common is RSI (*repetitive strain injury*), caused by long periods of typing or repetitive movements. RSI is an umbrella term for a variety of injuries which can be traced back to the use of computers in work environments. Specific terms include tenosynovitis, tendonitis, epicondylitis, carpal tunnel syndrome, and thoracic outlet syndrome. In 1997, a British judge dismissed the concept of RSI had 'no place in medical textbooks'. However, also in 1993, the use of VDUs in the workplace was introduced, in the form of the *Display Screen*



RSI is now a recognised problem, to the point where manufacturers put health warnings on their products. It is more of a problem than traditional typesetting, due to the lack of diversity of movement. A traditional typewriter required a lot of push of paper, and push the carriage. These are not the case with a computer, and it is perfectly possible for someone to type for hours more without a break. In addition, they are often harder on a typewriter than on a keyboard, and therefore not as likely to cause RSI as keyboards.

Other health problems include *eye strain*, caused by staring at a screen for a long period of time. This can be caused by sitting without moving in an uncomfortable or badly designed chair. The use of computers can be damaging to the individual's health. The directive on the use of computers states that:

- An employee's work must be planned so that work involving the prolonged use of computers is periodically interrupted by other activities.
- The workstations must be made of separate components (chair, keyboard, mouse) so that employees can adjust the station to suit them. Monitors must meet safety standards for glare screens. An adjustable chair should be provided.
- Employers must analyse the risks of physical and mental stress caused by the use of computers and take appropriate action to remedy the situation, including giving employees regular breaks from keyboard work.

You can now buy both software and hardware to help with these problems. Hardware includes ergonomic mice and keyboards, and accessories such as foot rests and document holders. Speech recognition and other software can help to automate tasks and therefore reduce the amount of typing needed.

As individuals we have the responsibility to ensure that the information, software, and hardware we use are not used maliciously, and to abide by the laws. When using computers we must consider how our actions will affect others.

For example, consider the possible implication of sharing personal or financial information. If someone is able to access this information and what could they gain from doing so?

**COPYRIGHT
PROTECTED**



The Ten Commandments of Computer Ethics

These were created in 1992 by the Computer Ethics Institute and form a guide for all computer and ICT organisations. They are fairly simplistic and follow common sense of good base to consider when developing software or using computers.

1. Do not use a computer to harm others.
2. Do not interfere with other people's computer work.
3. Do not snoop around in other people's computer files.
4. Do not use a computer to steal.
5. Do not use a computer to bear false witness.
6. Do not copy or use proprietary software for which you have not paid.
7. Do not use other people's computer resources without authorisation.
8. Do not appropriate other people's intellectual output.
9. Do think about the social consequences of the program you are writing.
10. Do always use a computer in ways that ensure consideration and respect for others.

Malicious Software

It is illegal to write and/or deliberately distribute malicious software such as worms and viruses. Most computer professionals find the idea of producing such software morally wrong. Programs such as viruses can cause serious problems to computer systems and are also 'out of your control'. The use of such programs can have severe consequences; for example, if a hospital computer system was attacked by malicious software, information and treatments could be seriously affected.

SOCIAL CONSEQUENCES

The social implications of computer use are ones that few consider. There is the Internet can create what is known as a 'two-tier' society, comprised of those who have access to the fast Internet, known as *information rich*, and those who cannot, who are *information poor*. This can be seen at an increasing rate in situations such as education where pupils are required to submit digital versions of their work, or work that requires research that would otherwise be done using traditional methods of research.

With the widespread use of the Internet, particularly among the younger generation, socialising with social peers has become much easier. This has meant that social relationships can now have very different geographical locations, so now people can socialise in groups that are defined strongly by interest or opinion rather than being limited by where they live.

The Internet and the reduction in the cost of electronics have also had an impact on the way we experience and discover entertainment material. These days, websites such as YouTube are signed to a music label, as musicians are able to publish their work free of charge. Musicians can publish their works for free to a large audience. Recording equipment has become so cheap that scale recording that at one time would have only been possible in expensive studios can now be performed in bedrooms and garages.

The amount of personal information that is accessible to businesses and collected by them has increased exponentially. These days, large businesses such as supermarkets track our shopping habits and preferences in order to target marketing and even shop layout to increase sales and make more money. Even the shops that we use today are only possible because of stock ordering and POS (point-of-sale) terminals allows such large-scale shops to operate and locally run and sourced shops out of business.

**COPYRIGHT
PROTECTED**



The social make-up of cities has begun to change because large industrialised areas were created during the Industrial Revolution and beyond. People have been looking for a better life in the countryside ever since. Computers have helped to make this possible, although it is not as wealthy within society. This has left a disproportionately large number of people in cities.

These aforementioned developments in digital technologies have drastically changed the way we live, made, and the way information is accessed and preserved has enabled the creation of large groups of people with very little effort. By monitoring data flows you can track personal data without the person's knowledge. Once this data has been gathered, it can be used to make money by selling it, publishing it or using it maliciously.

LEGAL CONSEQUENCES


Lawful computer use is no different to living by the laws of a country; there are rules that are illegal. With the growth of the Internet and integration into devices an enormous amount of data is being generated. There are many parallels that can be drawn from laws that are found in computer use that are similar to those that govern computer use and many of these are to protect the right to privacy and protection of data.

One industry that has been seriously affected by copyright issues is the music industry. Music files are copyrighted; it is illegal to broadcast or receive copies of them. However, services like iTunes and Napster made the illegal sharing of such files very straightforward and caters for a large number of users.

This is because there is so much Internet traffic that monitoring is an almost impossible task. Only a very small proportion of Internet pirates are caught. This means that in order to catch them, more methods have to be employed. These methods hope to limit access, ability to copy files, and once purchased. These have always been a concern, but since the widespread use of the Internet, they degrade each time they are copied and their distribution is fast and straightforward.

Modern approaches include:

- Spotify (shown on the right) allows users to stream music to their devices for a monthly fee. It also offers a free service, which has limitations such as lower streaming quality and the inclusion of advertisements.
- Systems such as BBC iPlayer use a Microsoft Digital Rights Management (DRM) system that makes sure that BBC programs can be downloaded and stored for a certain length of time, but they then expire, hence making it impossible to copy and broadcast them indefinitely.
- Windows media player formats (WMA) place a copy restriction on files so that they cannot be copied. However, they have been imported from other formats, however, this technique is easily bypassed.



Digital rights management helps make sure that artists, etc. are paid for the copyright that they own. It is argued that the technology provided by DRM should they change player or the company go bust, then the copyright is lost. It makes competition between providers difficult and stops a collaborative effort to create new content since the increased use of digital media.

With the nature of data technology what it is, the legislation governing its use is a challenge. One of the biggest challenges facing legislators is the act of copyright infringement. Before digital technology, the act of copyright was still difficult, but digital technology has made it significantly easier. The price of copying and distributing has fallen to zero. Music and software are the most affected. The challenge changes as people cannot see why they should continue to pay for these services.

**COPYRIGHT
PROTECTED**



Identity theft / fraud

In 2013 alone there were 13.1 million reported cases of identity theft, and it is expected that this will continue to rise in the foreseeable future. Techniques such as 'phishing' and 'spoofing' are common attacks seen on the Internet when someone is trying to gather information. A phishing attack is an attempt to gather personal information by masquerading as a reliable source. On the other hand, pharming is a social engineering method of diverting data flows.

Cybercrimes leave no *physical* evidence, and a well-performed data-gathering cyber-attack will leave little trace that an attack ever happened. To add to this, the Internet allows for near perfect anonymity for those that know how to protect themselves. Given that there is no police force dedicated to cyber-crime prevention and that an attack can be launched from any computer, it is next to impossible to intercept an attack while it is happening. Once the attack is over, very little can be found out about it because of the lack of evidence; tracing of an IP address can be checked but with the use of virtual private networks the IP will usually be a dead end.

CULTURAL ISSUES

When computers became more prevalent it became seen as a basic right to have access to technology. This has revolutionised teaching. Traditionally, teachers were the only source of knowledge, but computers have slowly replaced teachers as the main source of information. This presents a challenge: what about those who are in poverty or are less financially secure? Computers have not been bought by most schools to provide for those students that are less fortunate. Computers are expensive to buy and maintain, and are often comprised of outdated hardware to drastic effect.

It is the responsibility of the owner of a website to ensure that the data found on the website is of high quality; in real-world applications (i.e. a newspaper) this is easy to govern. On the Internet, however, it is difficult to complain, and it is governed by various acts that reduce libel printings and so on. Information on the computer/Internet use; information is largely unadulterated, but there are no safeguards in place to prevent dedicated to personal views which are promoted as fact that can be seen as propaganda. It is for this reason that the *quality of information* needs to be taken into account when using information from a reliable source.

It can also be said that cultural issues take into consideration how data can be used. Consider a research group that wanted to document the location of every small village in Africa; how will that information be used and how will that use affect the villages? Will there be any safeguards put in to prevent people using the information for their own purposes?

Speech synthesisers and voice recognition

Such programs as word processors, spreadsheets and databases can be used to create documents for display. They can combine voice recognition software, allowing the user to interact with the program using a speech synthesiser to provide a response to the user. Currently under development, this technology will allow blind and visually impaired people to surf the Internet.

A synthesiser must be developed with a number of factors in mind, including

- Compatibility with the desired software
- The quality of the synthesised voice
- The ability to pronounce Russian and other foreign texts

**COPYRIGHT
PROTECTED**



Braille terminals and printers

A Braille terminal consists of an array of 20–80 Braille cells, connected to the computer. This displays part of the screen display in Braille, so that a visually impaired person can read it.

Such terminals are very expensive, but are preferable to the voice-operated/synthesiser system in some contexts. It is highly suitable for:

- Programming
- Dealing with formatted texts
- Non-natural-language texts

There are a number of Braille printers on the market, allowing documents to be printed well in conjunction with the Braille interface. A computer Braille printer will take

- The operation of the printer by a visually impaired person
- The facility to print simultaneously in normal type
- The facility to print in Cyrillic Braille code
- The facility to print in Braille

Computer technology can be used to improve quality of life for disabled people. People with movement impairment can control equipment attached to a computer by using just a finger. A device has been developed to assist motion-impaired people. *IRVIS (Interactive Robotic Visual Interface)* system being developed to allow disabled users to interact with a GUI to facilitate control of a robot.

Information overload

There are many information systems available to executives and managers. They provide informed decisions, based on the most up-to-date information. Systems include intranets and email. However, one argument is that increased access to information is detrimental to business interest. A recent survey commissioned by Reuters for executive managers claimed to have been made ill by stress.

One of the key reasons cited was the overwhelming volume of information being available on the Internet named as a major source. The fact that so much information is available where they have to keep up, or else they will be disadvantaged. Though there is no doubt on this subject, it is clear from the Reuters survey that in some situations the quantity of information is a problem.

Computing can do more to help situations such as this, through more intelligent data processing. A solution to this problem lies not in reducing the information load, but in better organisation of the data. Both of these tasks can be performed by computers. A person would take a person to organise the same amount of data.

Forged documentation

Advances in computing and printing technology, combined with its increased affordability, has made it more possible to create accurately-looking forged documents. This has seen an increase in identity theft and illegal immigration. This kind of trend is a major impact of society and the level of control over identity is important to maintain a safe and secure society. Producing counterfeit money is another example of a problem that has become more widespread due to the level of accuracy that is achievable, due to technology.

In order to combat this, documents such as passports are becoming increasingly complex in terms of the security features they contain. For example, the most recent UK passport (2015) is printed using UV and infrared light, inks and watermarks and uses a single sheet of paper for the personal details page adjoined to the back cover to prevent it being tampered with. However, with the illegal trade of ID being so widespread, it will surely only be a matter of time until a method of duplication comes to light.

**COPYRIGHT
PROTECTED**



Questions: Computing Issues

- 1 You work for an organisation that would like to gather information regarding which calls are handled by employees. Discuss the potential ethical issues raised.
- 2 Consider the implications of pirating software and how the act of reproduction affects those who have produced it, those who are using it and those who have been affected.
- 3 What are the implications of a government flying reconnaissance drones over a sensitive area?
- 4 Research the health risk factors involved with heavy and prolonged use of computers and compare these with common health risks found 50 years ago.



INSPECTION COPY



INSPECTION COPY

INSPECTION COPY

**COPYRIGHT
PROTECTED**



9. Communication and Networking

The need for machines to communicate with each other has become an essential part of our lives. This section shows how we have developed our standards and methods to include networking. This, in turn, has revolutionised the world and its business and social links.

This section covers:

9.1 Communication	p1	9.3 The Internet	
9.2 Networking	p4	9.4 Transmission Control Protocol	

9.1 COMMUNICATION


COMMUNICATION METHODS

In order for a computer system to be responsive and useful it needs to have a user, its own input/output peripheral devices and its internal devices, so there are several methods of data transmission in a system that enable the computer to do just this.


Serial transmission

In serial transmission, the binary signals representing the data are transmitted over a single channel. There are several different classes of serial transmission:


- *Simplex* method allows communication in one direction only. This one-way communication is used for devices that only send or only receive data.



- *Half-duplex* method allows communication in both directions, but not at the same time. There is a single channel and the direction is switched after completing one direction.



- *Duplex method* allows communication in both directions at the same time. Two channels are permanently available. This is a necessity in interactive systems (where the user continuously requires a response).



Serial transmission can operate either in synchronous mode (with a mutual clock) or asynchronous mode (without a mutual clock). USB, SATA and RS232 are all examples of common serial communication standards, such as SATA which is used to connect hard drives. USB (universal serial bus) is a very common connection type used to connect various devices to a computer.

INSPECTION COPY

COPYRIGHT
PROTECTED



Parallel transmission

With parallel transmission more than one series of data bits can be transmitted over multiple cables in place of the single cable found in serial communication systems. Parallel transmission can be synchronous since there is a central clock that controls the timing of data transmission.

Parallel communication is only practical over very short distances due to the fact that as the wires get longer, parallel transmission suffers from skew. As the distance is increased, the lines move slightly out of sync and this massively reduces reliability.

Parallel communication used to be very commonly used to connect peripheral devices. However, due to the expense and limitations of the parallel system, serial communication is now commonly used. One place parallel communications can still be found is in the PCI bus, for example, is a parallel system. However, these are being replaced by serial communication. For example, PCI express is actually primarily a serial system.

Asynchronous data transmission

Asynchronous data transmission is a method of transferring data that has no fixed timing relationship between the order of the data and the time of the transmission. This means the transmitter's and receiver's clocks are not synchronised and only become synchronised at the time of the transmission.

Start and stop bits are used as markers to indicate where the message starts and ends. Start bits are added to them before sending. Start and stop bits are used to synchronise the two devices. Both devices must receive and send signals at the same rate. This is the reason why start and stop bits are required.

For example, the bit string 00101111001 would have the stop bit 0 and a start bit 1. So the transmitted string would be 0 00101111001 1.

COMMUNICATION BASICS

Before you can understand how data flows through a system and how components interact, you need to learn some basic principles and terms that you must learn.

Baud rate

Baud rate is the rate at which the signal changes (per second) in the communication system. The measurement used for a baud rate is called a baud. For example, 2 kBd means 2,000 bauds per second.

Baud rate is often confused with bit rate. The main misconception is that one baud is the same as sending one bit. However this may not be true, since one signal change may represent more than one bit. If we take a voltage of a wire to determine the bit rate, we can see that different voltages to signify more than one bit. For example, we could detect between 0V and 7.5V as more than one bit at a time, as below.

Signal Voltage Level	Binary Signal	Signal Voltage Level
0 V	0	0 V
2.5 V	1	2.5 V
5.0 V		5.0 V
7.5 V		7.5 V

Bit rate

Bit rate is the number of bits that can be transmitted per second. The units of bit rate are kbps (kilobits per second), Mbps (megabits per second) or Gbps (gigabits per second). For example, 9,600 kbps means that 9,600 bits of data can be transferred per second. A bit can be represented by a voltage level.

Bit rate = baud rate × bits per baud

COPYRIGHT
PROTECTED



Bandwidth

Bandwidth is the useful bit rate, also known as the information rate.

Usually, when transferring data between networks, some bits are required to handle communications which are considered 'wasted' bits. Bandwidth is also measured in bits per second. A bandwidth of 2 Mbps (2 megabits per second) indicates that there is 2 megabits of information lying in the communication medium. This means that the bit rate and the bandwidth are directly proportional to each other.

Latency

Latency is the time required for one bit to be transferred from one end to another between the sender sending a message and the receiver receiving the message in milliseconds. It is usually related to the length of the connection and the speed of routing/forwarding taking place.



Protocol

When data is being transmitted in a computer system, a set of rules that constrain how the transfer of data will be done. A set of rules established before transmission can take place is known as a protocol. Protocols are most commonly found in web-based applications and can be indexed.

Questions: Communication Methods

- 1 What is the difference between synchronous and asynchronous transmission? (1 mark)
- 2 Are the following true or false?
 - a) Baud rate is the same as bit rate. (1 mark)
 - b) Bandwidth is the time taken to receive packets of data. (1 mark)
 - c) Latency is the time taken for 1 bit to be transferred across a given medium. (1 mark)
- 3 Using what you know of the different buses, what controls the direction of transmission? (1 mark)



**COPYRIGHT
PROTECTED**



9.2 NETWORKING

NETWORK TOPOLOGY

Bus network topology

In this system, a single cable called the bus is used, to which the workstation is connected. This system is simple to run and normally uses less cable than any other system as one cable is used for the bus, and computers simply connect into it using short cables. However, if there is a break anywhere on the bus cable, it results in all the workstations being unable to communicate. Another disadvantage is that there is a larger degradation of performance under load than with some other topologies.

Advantages	
<ul style="list-style-type: none">■ Inexpensive■ Scalable – easy to add new nodes	<ul style="list-style-type: none">■ Single point of failure■ Network performance degrades

Star network topology

A star network has one central message-switching device through which all computers communicate. In this system, each workstation is connected directly to the central device – normally a switch – by its own unique cable. A major disadvantage, of course, is dependency on the central device, as all communications will cease if it breaks down.

The main advantage of the star topology is that it is easy to track down problems because each workstation has its own separate cable. Also, this means that a problem with one cable or workstation does not affect the rest of the system. Performance doesn't degrade when new nodes are added so the system is very scalable. A disadvantage is more cable to install a star topology, and the switch can be quite expensive. If a workstation crashes instead of a switch; however, if the computer crashes, then the other computers can still communicate with each other.

The star topology is now the most popular choice when a new network is installed.

Advantages	
<ul style="list-style-type: none">■ No performance degradation■ Secure communication – communication is direct therefore no eavesdropping is present■ Scalable – easy to add new nodes	<ul style="list-style-type: none">■ Expensive to implement■ Single point of failure – if the switch fails, communication is lost

INSPECTION COPY

COPYRIGHT
PROTECTED



Peer-to-peer networking

Peer-to-peer networks are those where there is no central server and every computer has the same rights. There is no central management. Peer-to-peer is useful when most of the time computers do not need to interchange information with each other. Small companies benefit from peer-to-peer networks since having independent computers does not require an expensive server or hardware to manage and maintain.

Programs are installed on every computer; however, careful management is required to ensure all software and data is up to date and protected. Sometimes a peer-to-peer network may have shared directories to store information on one machine; however, on most peer-to-peer networks the data is stored locally on each machine and so therefore each machine requires regular backups to prevent data loss.

Small networks often are peer-to-peer as the cost of maintenance does not vary with the number of machines. However, as the number of machines increases in an organisation, networks will justify the extra expense.

In February 2009, peer-to-peer networks have been estimated to collectively account for 70% of all Internet traffic (depending on geographical location). *BitTorrent* is a protocol for sharing files on the internet, and uses peer-to-peer networking to distribute data over the Internet. This involves multiple computers (or *peers*) using a BitTorrent client to communicate with each other. No central authority is needed to coordinate the download of files. The file is segmented into pieces and finding separate nodes to download those pieces. In 2013, BitTorrent was responsible for 3.35% of all worldwide bandwidth; more than 10% of bandwidth dedicated to file sharing.

Server-based networking

Also known as client-server networking, server-based networking is where one computer manages resources, security and other services. These services are provided by the server. The server is the only computer that holds all the information for the organisation; this is advantageous where a lot of interaction is required between computers. The organisation is able to control the access to the information better, provide better security, and use resources more efficiently (i.e. remotely) across a network rather than having each computer do its own thing.

Server-based networks are used in companies where people need to access resources all the time. A server can cope with many people connecting and gaining a resource without a drop in performance.

Security is centralised and users can authenticate themselves by using a single password to access the resources they are allowed to. This allows a user to be able to access their information on any computer in the network rather than relying on a single machine. There is also no need to worry about backing up data, since it is all done by the server itself. A lot of people can be added to the network without a drop in performance. Increasingly, home servers are also beginning to appear as the need to share data and resources around homes with multiple computers grows.

**COPYRIGHT
PROTECTED**



Wi-Fi

Wi-Fi is a high-bandwidth wireless method of communication which can be used in place of, or in combination with, wired Ethernet networks. Wired Ethernet networks are, however, still capable of far higher transfer rates and are more reliable than Wi-Fi networks. Wi-Fi networks are accessed through *hotspots* which are the areas within which a device can connect to a local area network, and are used to allow a device to connect to the Internet.

All Wi-Fi networks are based on an international standard so that, given your device can connect to any network without hassle.

Connecting to a wireless network

Wi-Fi hotspots are actually the broadcast area of what is known as the *wireless network*. A router is connected to the network itself. The device you need to be connected to the router directly using cables.

The other device you need to connect to a wireless network is either a *wireless network interface card* (attached externally to a computer via USB), or a *wireless network interface card* (a hardware component attached directly to the motherboard (in desktop computers) or the PCI slots). Both of these act as an interpreter – to send and receive data that is easily transmitted.

The type of network adapter needed will vary depending on the protocols it supports, the communication medium and the topology of the network to be connected to.

How networks are secured

There are a number of ways in which you can improve the security on a wireless network. Over time, a number of protocols have been developed over time to help secure a network and prevent *rogue users* being able to access the network.

The methods that you are required to know are:

- *Wi-Fi Protected Access (WPA)* versions 1 and 2
- *Service Set Identification (SSID)* broadcasting
- MAC address whitelisting

Wi-Fi Protected Access (WPA/WPA2)

WPA was introduced in 2003 to address the shortcomings of Wireless Encryption Protocol (WEP). When you connect to a network that uses WPA as a security procedure you'll need to enter in a passphrase; this acts as a layer of security. If you cannot connect to the network, however, when you connect to the network, the router communicates and creates an encryption key. The key is generated using a random key technique for each device connected to the network. Security generated keys are temporal. Temporal key integrity protocol (TKIP) is a protocol that checks to ensure that the keys have not been tampered with; when the encryption type between devices is changed.

INSPECTION COPY

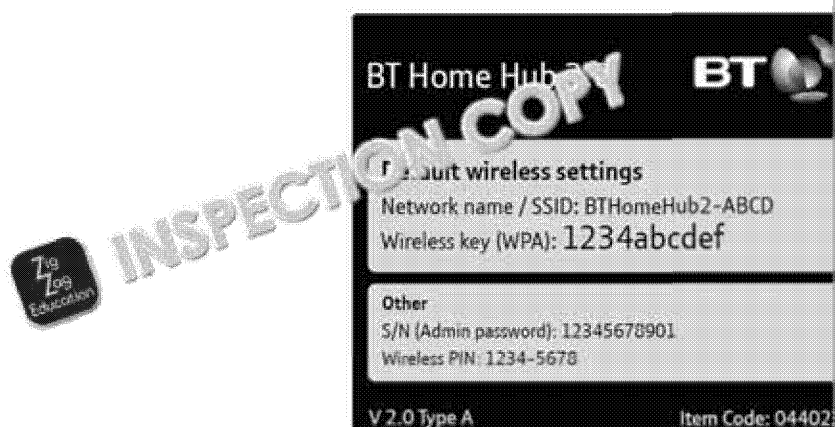
COPYRIGHT
PROTECTED



Service Set Identification (SSID)

An SSID is the name given to a wireless local area network by the administrator. Any device connected to a network must have the same SSID employed to their connect to each other through the network, otherwise they are invisible to each other.

Think about your home network; the odds are that your router broadcasts its SSID so that if the device knows the WEP key it can connect. A way of improving this security would be to switch off the automatic SSID broadcast; that way it becomes invisible to anyone that doesn't know the *exact* name of your Wi-Fi can connect to it.



MAC address whitelisting

Many administrators of WLAN will use a MAC address whitelist to control who can access the network. The whitelist effectively acts as a spam filter to block every device not on the whitelist table, which can only be accessed by logging in to the router. This is because the data being sent over a network simply isn't feasible.

Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)

CSMA/CA is a transmission protocol used in networks that acts to prevent the collision of data between nodes. As soon as a node receives a packet that is to be sent on the network, it checks if the channel is clear. If there is no available channel, a back-off time is generated, at which point the node will check again.

If a packet of data that is larger than a predetermined size is needed to be sent, there must be a certain *handshake* that needs to happen; this is called the *Request to Send* (RTS) and *Clear to Send* (CTS). This protocol only takes effect if the packet is larger than the threshold because it will take more bandwidth to send than a smaller packet.

Questions: Networking

- 1 What are the advantages of a star network topology? (3 marks)
- 2 What piece of hardware is required to access a wireless network? (1 mark)
- 3 Explain how you could use a combination of networking methods to improve the security of a network. (6 marks)

**COPYRIGHT
PROTECTED**



9.3 THE INTERNET

THE INTERNET AND HOW IT WORKS

The structure of the Internet

The Internet is vast, and the Internet you see on a day-to-day basis when using the Internet is just the tip of the iceberg – but if the Internet is truly that large, what is it?

The hardware that makes up the Internet is made up of *servers* which hold Internet servers together and other equipment (*routers*) that link computers to the Internet. The software that makes up the Internet is invisible to people who use the Internet. It produces the flow of information from one side of the world to another, and organises the email system.

Packets

For data to be sent across a network it is broken down into a structure known as a packet. To give you some idea, a packet is about 100 bytes in size, so a 2 megabyte file will be broadcast as 20,000 packets. The network so the most direct route to the destination will change during the course of the journey, so packets will arrive in a different order to that which they were sent in.

To enable packets to reach the correct destinations and then be assembled into the original information is stored in part of the packet known as the *packet header*.

The header consists of:

- The destination address
- The source identification
- A checksum (for transmission error detection)
- The packet sequence number

Destination
Source
Checksum
Sequence

Packet switching

Packet switches have two main functions:

1. To enable more than one device to share a (usually high-speed) data link
2. To find the most direct route for information to travel

Packet switching was developed to replace circuit switching and to make the network more autonomous. For example, the telephone system originally used dedicated circuits for each conversation. This is clearly suboptimal, as a normal conversation would tie up a circuit for the duration of the call, which could use far more bandwidth. A more efficient system is packet switching, where data is sent across a network as they are required, which results in a potentially increased efficiency.

Routers

A *router* is a device that forwards packets from one network to another. Each network has a router which works out where to send packets which are destined for the network they are actually within. When a packet is sent across the Internet, it will go through many routers before it reaches its destination.

Routers are able to work because every piece of network hardware (with the exception of a hub) has a unique MAC address (see 9.2.3). Routers can therefore change the destination MAC address of a packet without changing the destination IP address. This means that routers can forward packets to another, changing the destination MAC address as they go, while maintaining the destination IP address. Each router on the path is therefore able to work out where next to send the packet to reach the destination IP address.

INSPECTION COPY

COPYRIGHT
PROTECTED



Suppose a user wants to send information from one computer to another computer. The user prepares the message and sends it off to the router. The router obtains the forwarding table. When the route is found, the message is transferred to another router. In its network the destination node then the router transfers the message to the destination. The message to another router which performs the same process. All routers are connected to each other, which makes communication universal, and are able to transfer data across different architectures.

Gateways

Gateways are the entrance and exit of a network and can be considered as the main use of a gateway is to connect multiple networks of different architectures. The gateway doesn't have to be a physical object, though; it can be implemented into the network software which can then be distributed throughout the network. The gateway is physically bound to a single location. The gateway takes in a packet from an external network, removing all the information apart from the raw data. The raw data is then re-encapsulated into the protocol which that network supports so that the data can be sent to where it appears in a network map, usually towards the edges and access points. The gateway also ensures the network security that has been employed by the network administrator.

Uniform Resource Locator (URL) and internetworking

An internetwork is a collection of independent networks connected to each other. Each network functions on its own and does not depend on other networks. An example of this is the Internet. The networks in an internetwork are joined together with routers which handle the URL requests to find certain volumes of information. The Uniform Resource Identifier can be either a Uniform Resource Locator (URL) or Uniform Resource Name (URN). A URI provides a unique reference for an Internet resource while not necessarily providing an exact location for the resource as it could be a query in a database or call to some application. This is why the URL is an identifying location for a resource providing a means of getting the resource.

For example, the URL **http://www.google.com** is the URI that exclusively identifies the Google homepage. It describes the type of representation that is implied that the Hypertext Transfer Protocol (HTTP) must be used to collect the data. On the other hand, provides a way of exclusively identifying something, for instance an ISBN number identifies a specific book, but provides no means of actually getting hold of the book. A phone number is equivalent to a person's phone number or postal address, while a URN is equivalent to a person's phone number or postal address.

Domain names and IP addresses

Domain names provide the text that is entered in the navigation bar of the web browser. A domain name service (explained below) to return the specific IP address of the domain name. For example, the address **www.google.com**, then **'www.google.com'** is the domain name which identifies the top-level part of the domain and is often related to the type of service provided. The **'google'** part of the domain, known as the second-level identifier for a particular IP address. The **'www'** provides the name of the host server that would deal with file transfer protocol (FTP) requests. Also, in the hierarchy between the secondary and server part of the address; for example, **'www.images.google.com'**.

**COPYRIGHT
PROTECTED**



Internet registries

If you stopped to think about IP addresses for a moment you'd soon realise their value, so an impartial system was required in order to facilitate the distribution of consequently IP addresses. This is the role of Internet registries and registrar organisation called the Internet Assigned Numbers Authority (IANA) that decides in which regions, and the local Internet Registries and their registrars that are assigned to them according to local legislation, etc. An example of a UK Internet registry is **www.nominet.org.uk**. If you had just founded a company it would be a registrar's approach in order to register the domain name of your new company. This of course is not free and you will also often have to 'rent' the web space where your web pages are stored so that they can be accessed by others.

These companies will also often be involved with legal proceedings when it is found that they should not own specific domain names. For example, when they are profiting from another company other than Ford's, as Ford had owned **www.ford.com** before it was a trademark, that they were profiting from Ford's name, then the domain name might be given back to Ford Motors.

INTERNET SECURITY

Firewalls

Firewalls are a very common and easy method of protecting your network, and subsequently your system, from the risks of using the Internet. Traditionally, users of a network would rely on the security of the individual hosts to protect them but as the number of hosts increases it made it less manageable; there was more likely going to be a lapse in security due to the drop in uniformity of security. Personal firewalls were developed to combat this and use services to provide the protection.

Packet-filtering firewall

Packet filtering operates at the packet level and studies each packet as it arrives and passes through router interfaces depending on the filtering protocol that has been put in place, which is checked against the features of the packets themselves.

Filtering protocols are put in place by the network administrator or someone who is overseeing the protection of the network. There are obviously too many packets involved in data transfer to monitor each packet that attempts to gain access to a network; this is why the network administrator will employ several policies to regulate the ports of the network.

Advantages	Disadvantages
<ul style="list-style-type: none"> ■ Firewall is provided by standard router ■ Fast ■ Flexible ■ No user action required for installation 	<ul style="list-style-type: none"> ■ Filtering rules are vulnerable ■ Cannot be tested for effectiveness ■ Routers may only be able to filter when a break-in occurs ■ May not be able to filter all traffic

COPYRIGHT
PROTECTED



Stateful inspection firewall

These firewalls have no built-in concept of 'state' or 'context' – that is, the firewall is for, where it is going or what it does. This type of firewall examines every packet and prevents any packets from entering or leaving a network if they do not meet the criteria. The process of packet examination is *very* resource-intensive as there is constantly a stream of packets being checked by the system at any given time. It stores the *states* of connections into a state table and uses hashing of the data to be processed more quickly; the states that are stored in the table are compared against the states of the allowed policies. However, if the first packet of a connection is standard then there is no need to check the following packets associated with it.

Advantages	
<ul style="list-style-type: none"> ■ Faster than using a proxy server ■ Due to its nature it has a higher level of security built in by performing application layer filtering of packets ■ More secure than basic packet filtering ■ Cheap to set up ■ Flexible but strict on the rule set 	<ul style="list-style-type: none"> ■ Can be less secure ■ Slower than basic packet filtering ■ Very resource-demanding ■ Can be vulnerable to attacks that target the protocol ■ Creating the policy set can be difficult

Proxy server

A proxy server is often used with a firewall to increase the security of the system. Computers inside a network use proxy servers to access external information. The main advantage of this is that only one computer is exposed to the outside network. The rules to be 'fine-tuned' to allow control over connections.

Advantages	
<ul style="list-style-type: none"> ■ Only allows proxy services through ■ Protocols can be filtered and manipulated ■ Hides internal structure through information hiding ■ Improved authentication and logging ■ Cost-effective ■ Rule set is less complex than packet filtering 	<ul style="list-style-type: none"> ■ Protocols are more secure ■ Authentication is easier ■ Editing protocols is easier ■ Cost-effective

**COPYRIGHT
PROTECTED**



SYMMETRIC AND ASYMMETRIC ENCRYPTION


Encryption is a technique used to protect data by making it unreadable. Spec used to convert the data, which is in plain text, to cipher text. Plain text is th is the original data that has been transformed into a completely unreadable t understand without the use of a key. Keys are used to decrypt messages into messages into cipher text. Symmetric encryption uses the same key for encry asymmetric uses different keys for encryption and decryption.

In public key cryptography, each party using a public/private key encryption s known only to them, and a public key, which is freely available. These keys a private key can decode messages encoded by the public key, and the public encrypted with the private key. Importantly, however, the public key cannot encoded by itself and the same goes for the private key. To send an encrypted transmitter will encode the mes, a using the public key of the receiver. On private key to decode the message and so only they can decode it.

Messa try ed using a public key can only be decrypted using the corre way, m es encrypted using the private key can only be decrypted using t what allows the process of digital signatures.

Symmetric encryption

As mentioned above, symmetric encryption is where the key that is used to s encrypt and decrypt a message. Technically, it is the inverse of the key that is the data required to complete a process is derived from the secret key that is feature of the application that is being used to transmit the data.



A good example of a symmetric encryption key that was used widely was the which dates back to 1970. When it was first introduced the developers said t symmetric key's 52-bit secret key, but it soon came under scrutiny because c be secure. In 1997 there was a joint effort of 14,000 as an attempt to try and and it took four months of continuous computation to derive the key's value. was built for £250,000 which managed to crack the algorithm in just three d

INSPECTION COPY

COPYRIGHT
PROTECTED



Asymmetric keys

The opposite of the symmetric key is the asymmetric key. These are frequently used in the business world because they offer one of the highest forms of protection. The

- The sender doesn't need to decrypt the data once they've sent it
- Only certain people should be able to read the encrypted data
- By separating the keys people can only send data


By separating the keys and making them different, it allows you to control who can read your messages because once the message or data is encrypted it becomes unreadable. The sender has a private key to go with it. This is because with the symmetric encryption you would have to use the same key to encrypt and decrypt the message. If the key were affected and still decrypt the message to retrieve the plain text.

Without the private key it becomes almost impossible to retrieve the plain text. This improves the integrity of the message being sent.

Did you know?!

There is a wide variety of security threats that have been developed by people and organisations. For example, one of the most prevalent threats that has been making its way around the world is ransomware. This is where a virus has infected a computer and instead of self-replicating and spreading to other devices on the network, it will encrypt the data on the hard drive with a key with which only the attacker can unlock. If you don't pay the ransom then they don't unlock your computer.

The public key is made available to everyone; it is either found on a company's website or given to them by other means. The sender then encrypts the file using the public key and sends it to the receiver. Malicious users can try to intercept the file, but if they do it is just a jumble of cipher text; other malicious users can still send traffic using the public key but it will be blocked by a spam filter on the receiver's end. The receiver can then decrypt the message using their private key to produce the plain-text representation.



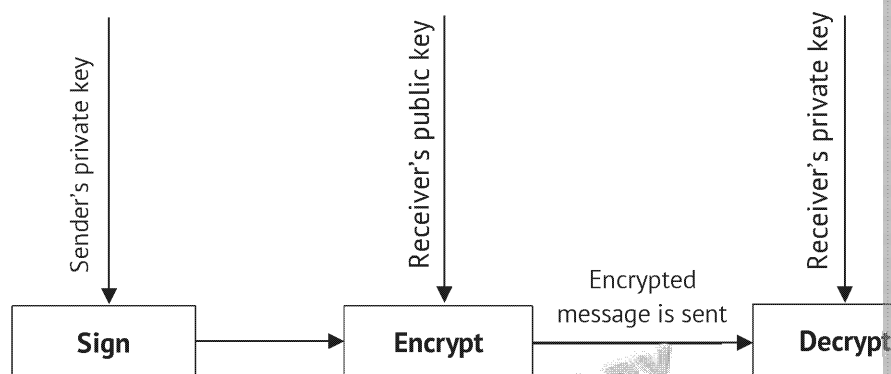
DIGITAL SIGNATURES

Digital signatures are used in conjunction with public key cryptography. Public key cryptography ensures the data is secure but it does not guarantee that the message being sent has not been tampered with. With a digital signature, the person who receives the message is able to tell who the sender is and whether the message has been tampered with.

A digital signature is generated by analysing the document and formalising it into a unique code. Using the mathematical notation, a hash function acts on the data and generates a unique code that depends on the content of the message. If any of the contents change then the hash function will produce a different code and show that it has been tampered with. Digital signatures are also used by programmers to sign their applications, which shows that a specific application

**COPYRIGHT
PROTECTED**





DIGITAL CERTIFICATES

A digital certificate is a type of ID card which lets you identify a specific user. It contains information such as a serial number, the user's public key and the certificate authority's signature. Certificates are used to verify that the sender is authentic. One reason for this is that maybe someone pretends to be someone else and sends you a message. You can encrypt the message using their public key instead of the intended person's. In this situation, a certificate authority gives you a certificate which verifies your identity. If someone wanted to send you some piece of data, then they should include their certificate. The certificate authority is contacted to verify that the certificate is valid.

SECURITY THREATS

Viruses

A *virus* is a computer program embedded into another apparently harmless program. It causes harm to a computer. The first step that a computer virus performs when it is executed is to copy itself onto disk and hide itself. After being copied onto disk, the virus can replicate itself on the system so that it causes problems. Viruses tend to create multiple copies of themselves and spread to other computers. Antivirus programs are used to detect and remove these viruses. Many operating systems have built-in virus scanners which scan files that are available for download.

An infamous example of a virus would be that of the Commwarrior, the first virus to spread via multimedia messages or via Bluetooth. The virus would be sent by an MMS; when the user received the message then the .SIS file would install itself onto the phone; it would then send out messages containing a copy of itself to other phones from the contact list. The virus spread through Bluetooth connectivity; when the phone was used the virus would attempt to infect other phones every minute and try to send a copy of itself to them.

Worms

A *worm* is a malicious program designed to replicate itself in an attempt to spread across a network. The most significant difference between a worm and a virus is that a worm does not need a host program to run. Worms can disrupt entire networks, causing traffic, denial of service, and data theft. The worms spread by several means but the most common is by email; although they cause damage to the host device until their payload has been delivered, they can also be used to control at which traffic is handled in a service.

Consider a server that handles the call requests for a business's login details. One of the possibilities is that every time a data call is made to the server, instead of the server returning the login details, a worm could return a copy of itself. The user wouldn't log in, so they would try again. They wouldn't be any the wiser that their computer had now been infected.

**COPYRIGHT
PROTECTED**



Did you know?!

One of the most famous worms that has been seen is that of the *Morris* worm. It was created by a Cornell University student in the USA. The student wanted to see how big the Internet was. It consisted of lines of code to trace an IP address, return the IP address and search the host computer for files when files were transferred. Although the code wasn't intended to be malicious, it was that made the host service very unstable and often resulted in the machine being taken offline.

Trojan

The term *Trojan* is derived from the Trojan horse used to defeat Troy by the Greeks. It is a replicating virus that is hidden in a downloaded file and is unleashed when it is run to gain access into a system. The Trojan program then acts as a back door for someone to access the infected system.

Not to leave Mac users out from potential threats found in modern computing, a virus was found circulating Apple computers in 2006. The virus was called 'Leap' and it would be transmitted over the Internet. Instead the virus would affect local area networks, such as home with multiple computers or those found in the business world. Leap was an instant chat program found for the Mac OS. The virus would then target other programs to gain access to the most secure parts of the computer. It would ask non-admin users for admin rights; the admin would then put their password in which would be logged and used to access the whole of the computer and would use the instant messenger's network to spread to other computers and send a copy of itself.

Did you know?!

One of the perceived advantages of Apple Mac computers is that there are no viruses. However, this perception isn't entirely accurate. In practice this is true because of the way an Apple machine is designed. The problem, or solution, with transmitting viruses on a Mac is that you can't run third-party applications without explicit permissions, so you can infect a Mac but it will be slow and limited. Apple products are no harder to infect than any other machine. The first virus ever produced on a Mac was by Rich Skrenta in 1982. Skrenta, a computer hacker, created a cloner which was capable of infecting the Mac's boot sector, a feat that is still impressive today. This virus predates the first IBM virus (the 'Brain' virus) by nearly five years.

Questions: The Internet

1. A packet is said to have a payload of 64 bits. What does this mean? (1 mark)
2. What is packet sniffing and why is it useful? (2 marks)
3. What is a packet-filtering firewall? (1 marks)
4. What is the difference between a computer virus and a Trojan? (2 marks)

**COPYRIGHT
PROTECTED**

9.4 TRANSMISSION CONTROL PROTOCOL / INTERNET PROTOCOL

TCP/IP

TCP/IP stands for *Transmission Control Protocol / Internet Protocol*. It is a protocol that allows communications on LAN and WAN networks and is widely considered the standard for Internet communications. The protocol is arranged in a stack with four layers which are:

1. Application Layer
2. Transport Layer
3. Network Layer (IP)
4. Link Layer

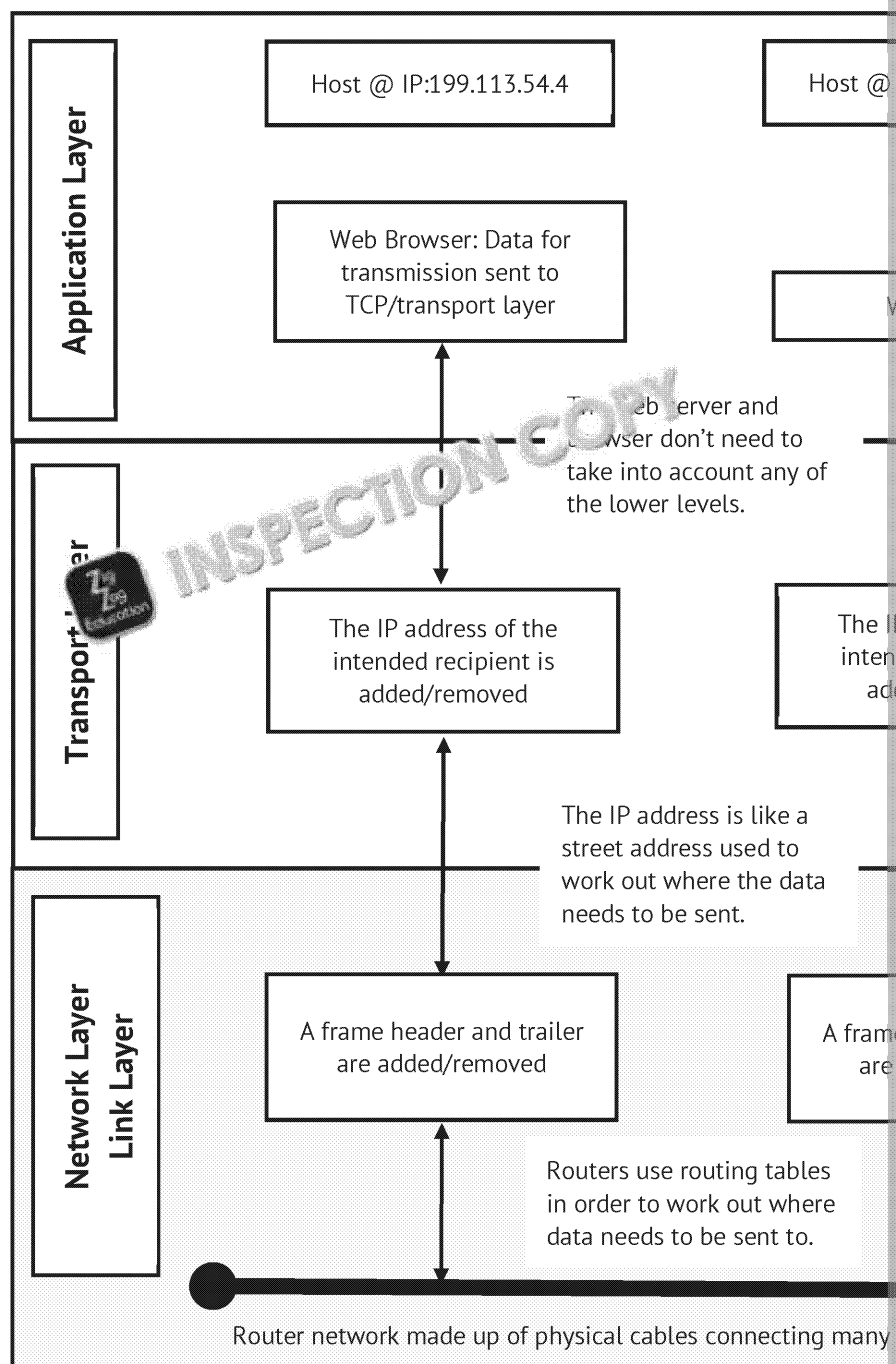
This layer processes data at the application layer, which is the actual data being transmitted. It goes through several stages before it can actually be transmitted. In the transport layer, the data is broken down into segments. In the network layer, the segments are given an IP address to distinguish it from other communications on the network. Once these stages are complete, the information is broadcast to the intended recipient, where the stages are reversed to retrieve the original data. This involves removing the header and trailer and also the IP address. The data is then reassembled into separate packets. The recipient can then process the application data as it was intended. Packets are labelled with an order position because often they will get disrupted and need reordering on receipt.

The complexities of each of these layers are removed by the introduction of a programming language that allows a programmer to have any involvement in the layers below the application layer. This makes the use of Internet applications much more straightforward. They often allow socket programming to file paths and then used to transmit data to the intended recipient. A socket is a unique identifier, such as `'www.google.com:8080'` which sends the request to the IP address associated with the port number 8080, which may be waiting to receive information from a given application. The data is then rendered into a web page in text format once a connection has been made. This will be used by the browser to render the Internet resources correctly and control the content being displayed. The following diagram shows how the different parts of the protocol fit together.

INSPECTION COPY

COPYRIGHT
PROTECTED





Well-known ports

Ports are the computer software way of exchanging data directly; they each relates to the kind of network it is. Often certain protocols and applications are assigned specific port numbers from 0 to 65535. The well-known ports and contain entries such as:

20 – FTP	80 – HTTP
21 – control FTP	110 – POP3
25 – Telnet	443 – HTTPS

The port numbers from 1024–49151 are the registered ports which have been assigned by the *Assigned Numbers Authority* (IANA). The port numbers from 49152 are the private ports. The port numbers outside of these three ranges are known as the ephemeral ports.

**COPYRIGHT
PROTECTED**



Media access control (MAC) address

MAC addresses are a unique 6-byte (48-bit) identifier that all network interfaces have in order to communicate with a network. The MAC address is assigned to a device by the manufacturer and, unlike an IP address of network mask, cannot be changed once it is assigned. It is used by the media access controller in the link layer of the TCP/IP model.

All MAC addresses on a network are kept in a table by the router for the network with an available subnet address so that the packets that you're trying to receive are not someone else's while also helping the router to maintain what addresses

STANDARD APPLICATION LAYER PROTOCOLS

Hypertext Transfer Protocol (HTTP)

This is the protocol that defines how data from web pages is transferred from the server to the client when the page is being viewed over a TCP/IP network. It is common to see the beginning of a web address, for example **http://www.google.com**. The HTTP protocol uses the IP address and the server will usually return a web page or the like to the client. There are eight HTTP commands that are combined to allow all the required functionality. GET requests a specific resource, POST which returns information, typically from a form. SSL is used among other things to create an HTTPS connection (which is discussed later). To make use of this protocol and to be able to retrieve web pages in text form, browsers like Java require Internet browser plugins which are initiated by these protocols.

File Transfer Protocol (FTP)

The File Transfer Protocol is what is used to download and upload files between computers.

Downloading is the term used for copying files from the Internet to your own computer. You will have noticed that Internet addresses have *http* at the beginning. This tells your Internet browser that you are downloading a web page from the Internet. Sometimes you want to download a different sort of file from the Internet, for example a trial program, a collection of pictures or a word-processing document. This is carried out using *FTP*.

There are also many FTP sites around the world. These are computers that hold other files that are freely available to everyone. Most major companies on the Internet (Microsoft, IBM and Novell). Most modern browsers have FTP built into them, so when you are *downloading* such files you won't need a special FTP program, although if you are downloading files it may be worth your while to try one.

If large files are being distributed, then it is quicker to collect them using FTP. It is reliable because FTP connections can often continue to download if something happens to the connection rather than starting from scratch. Uploading means transferring files onto the Internet. You would do this if you have your own Internet site and you want to write to your web space on the Internet. Again *FTP* is used, and most modern browsers have upload functionality built in as long as you have the correct permissions for it.

**COPYRIGHT
PROTECTED**



Secure Hypertext Transfer Protocol (HTTPS)

The HTTPS scheme is very similar to HTTP except that it connects to a different standard port and there is a layer of encryption between HTTP and TCP protocols. In standard HTTP the information that is transmitted could be intercepted and decoded by anyone, which makes the information transmitted over those connections insecure. In situations such as ecommerce and e-banking this level of security is not high enough.

The HTTPS protocol encrypts the data before it is transmitted over the network and sends it to a different port to the standard HTTP protocol. This is suitable for private information.

Post Office Protocol version 3 (POP3)

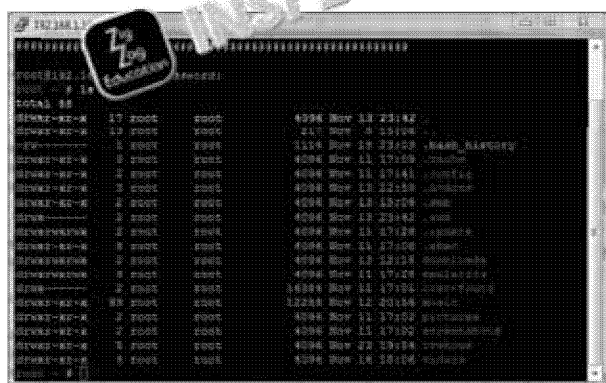
This protocol allows an email client to receive emails from an email server. The email server, however, does not delete the messages, save them as new messages on the user's mail client. This means that once the messages have been accessed, they are no longer available on the server. A newer protocol, IMAP (Internet Message Access Protocol), is suitable for email users who have two modes of email viewing, one of which is for messages to be downloaded and saved, but also from a web browser where messages are accessed through an online application. In this protocol the messages are only deleted when the user explicitly requests that they are. The main difference between these two protocols is that POP3 uses the user's mail store as the primary store and IMAP uses the server's mail store as the primary store. The email server will be located with the ISP or at the location of the email service company being used. This is the location that email is sent to and where they are received from using the POP3 and SMTP protocols.

Simple Mail Transfer Protocol (SMTP)

SMTP stands for Simple Mail Transfer Protocol and is used for the distribution of email. Each message will include the intended recipients and the message as well as other information that is necessary. These are then sent to the sender's mail server which will use the correct mail server for the recipient, who will use a POP3 or IMAP protocol to retrieve the message.

Secure Shell (SSH)

Secure Shell is a remote-access protocol that allows secure communication between a client and a server. The user has to be using the SSH client while the server must be running the SSH daemon. The connection is made over a secure channel on a potentially insecure network. The user can see the data transferred. It is used to log in to a remote server and execute commands across the server. The connection is made using TCP to a port of 22. It is also possible to use other application-level protocols to perform certain tasks. For example, to access a company's mail server the POP3 or IMAP protocols can be used to retrieve email, or the curl command for HTML.



This image shows an example of how to access the directory of a remote server using the SSH client. The terminal window displays the output of the 'ls' command, showing a list of files and folders. The text 'Embedded Linux Enterprise Edition' and 'a Windows network' are visible in the background of the terminal window.

**COPYRIGHT
PROTECTED**



IP ADDRESS STRUCTURE

An IP address should be a value unique to a device so that packets can be co between devices, otherwise an *IP conflict* will occur as packets are distributed the same IP address. All IP addresses are made up of four quadrants separated by dots, representing a binary address:

XXX · XXX · XXX · XXX

IP address: 192 · 168 · 10 · 5

Binary: 11000000 · 10101000 · 00001010 · 00000101

If you liken an IP address to your home address you can understand better how to identify devices. Your home address will have a county part and a city part. A network will have the same 'county' part and the devices will be given separate addresses.

For example, say there is a network with three devices on it: a server, a printer and a computer.

Server: 10.30.15.1	Computer: 10.30.15.5	Printer: 10.30.15.10
--------------------	----------------------	----------------------

You can see that the company's IP is in the first three quadrants and is '10.30.15'. This is called the company's *host address*.

IP class

There are also different classes of IP which were introduced in the *classful* Internet. The class system was developed to assist in the distribution of IP addresses and is shown at the table below.

Address class	First octets of address		Range of first octet values
Class A	0xxx	xxxx	1 to 126
Class B	10xx	xxxx	128 to 191
Class C	110x	xxxx	192 to 223

Class A	Class B
126 routable addresses For example: 84.42.199.5	16,384 routable addresses For example: 140.153.82.254

SUBNET MASKING

A subnet mask tells a computer which IP addresses it is able to reach directly (on the same switch) and which it cannot reach directly and therefore needs to access through a router. Subnet masks are in the same form as IP addresses; however, each number is either 255 or 0. 255 means that that part of the IP address is used to identify the network. A 0 means that that part of the IP address can still be used directly. The mask itself is a 32-bit number that is applied using a bit-wise AND comparison to split the address into the network identifier and the host identifier. The result is only true if *both* the values are the same and is set to 0 if they are not.

Address class	Bits for subnet mask			
Class A	11111111	00000000	00000000	00000000
Class B	11111111	11111111	00000000	00000000
Class C	11111111	11111111	11111111	00000000

**COPYRIGHT
PROTECTED**



Therefore, an example of a B class network with the IP address 138.96.0.0 with network notation would be written as 138.96.0.0/16. Take a look at the following table to see how this is used.

In this worked example you will see how to find the network address for the IP address 129.56.189.41 with the subnet mask 255.255.240.0.

10000001	00111000	10111101	00101001	IP address
11111111	11111111	11110000	00000000	Subnet mask
10000001	00111000	10110000	00000000	Network address

The IP address 129.56.189.41 therefore has a network address of 129.56.160.0 with a usable host identifier of 129.56.160.1 to 129.56.160.254.

IP Standards

There are currently two standards of IP address that are in use today. These are the Internet Protocol and are *IPv4* and *IPv6*. These are the two fundamental technologies that devices need to connect to the Internet.

IPv4

IPv4 is the 32-bit IP address that everyone is familiar with and that has been discussed on the previous page for explaining IP addresses. The 32-bit IPv4 IP address means there are 4.29 billion addresses available (roughly 4.29 billion addresses). However, since its inception there has been a prediction and in 2011 it was noted that the last of the IPv4 addresses had been used, though; the majority of addresses are dormant addresses owned by conglomerate companies, but it does mean that in the near future we will need a new standard (IPv6). An example of an IPv4 address would be 192.168.0.0, a Class C address.

IPv6

The move to IPv6 is inevitable; IPv4 addresses are becoming a dwindling resource. IPv6 is similar to its predecessor in that it assigns a unique number to each device that connects and communicates with a network, but it has one major difference. It uses a 128-bit address which will vastly increase the pool of available IP addresses. IPv4 has a pool of 2^{32} addresses, whereas 2^{128} grants an IP address pool of 3.40×10^{38} – three hundred and forty billion billion addresses. There are numerous advantages of IPv6 over IPv4. The large increase in address means that instead of writing the IP address in four groups of three octets you would for IPv4 you would write it in eight groups of two hexadecimals.

An example would be `4aae:1803:1c4f::2:0:b9be:de12:94cd`.

Did you know?!

Are you struggling to get your head round the size of the IP pool for IPv6? Let's put it into perspective. If you were to generate one million IP addresses per day, every day of the year it would take 9.32×10^{29} years to exhaust the pool. Another way 6.7×10^{19} times the current age of the universe. Try to write these numbers in a way that makes sense. Compare these to the number of seconds in a year.

**COPYRIGHT
PROTECTED**



Advantages of IPv6

- Removes the need for *Network Address Translation (NAT)* as all devices are given a public IP address
- By removing NAT you remove the possibility of private address collisions
- Easier to administrate – naturally leads to the removal of *Dynamic Host Configuration Protocol (DHCP)*
- Removes the need for a system administrator to set up complex network configurations
- Higher data transmission speed due to direct routing of packets
- Increase in security due to packet routing and the size of the IP pool
- Network infrastructure becomes simplified

Disadvantages of IPv6

- Specifying IP addresses becomes considerably harder with the increased size of the address space
- Internet service providers are given absolute control over their network, denial of service or restricted service by throttling an IP range
- There has been no clear decision as to when or how the Internet protocol transition and how this will affect the service they are providing to the user
- The transition will be very complex for everyone involved as the move is a phased transition to get the core components functional
- All machines that operate on IPv4 will become obsolete almost immediately which means the transition will be costly for everyone
- Potential rise in the cost of computers – with everyone needing new Internet service, manufacturers could charge whatever they wanted for it

PUBLIC AND PRIVATE IP ADDRESSES**Routable IP addresses**

Public IP addresses are assigned by IANA (Internet Assigned Numbers Authority) to localised registrars, who then assign numbers to individual users/companies. Every address on the Internet is unique, otherwise routers would not know where to send a packet. The public IP address of your router, not your computer. When you connect to the Internet, your computer is the middle man for all the packet traffic going in and out of your home network. Routers are used as a way of controlling the number of applied IP addresses. If every computer had a public IP address then we would have run out of IP addresses several years ago, but with the use of routers the distribution of packets is taken over by a network router.

Non-routable IP addresses

IANA has specified certain IP addresses as private. These are for local networks and are not connected to the Internet. To connect a LAN using these IP addresses to the Internet a *Network Address Translation (NAT)* must be used. All routers universally use NAT to work on private IP addresses. It is supposed to end up with the router changing outgoing packets so that they appear to be from a public IP address.

An analogy for this would be the internal post at a company. The Royal Mail takes the internal office block (the public IP address) and then the internal post works out which office it is (the private IP address).

Private IP addresses fall within the following ranges:

- 10.0.0.0–10.255.255.255
- 172.16.0.0–172.31.255.255
- 192.168.0.0–192.168.255.255

**COPYRIGHT
PROTECTED**



DYNAMIC HOST CONFIGURATION PROTOCOL (DHCP)


Dynamic Host Configuration is the protocol that supplies an IP host with all the data across a network. Without this protocol, if you were to move a computer to the network admin would have to manually recalibrate the network so that it can operate in the new subnet location, the IP address of the computer from the old location would be reclaimed manually and the process would be lengthy.

With the protocol, however, the entire process is automated. The protocol manages IP addresses on a network for devices and it operates on what is called a *lease* period. When a device joins a network it sends a request to the DHCP server which will service the request by providing available internal IP addresses that can be assigned to the device and allocate a policy that allows the device to connect to the network until it expires because of an elapsed time frame. When a lease expires, the IP address is added back into the pool of available addresses.

NETWORK ADDRESS TRANSLATION (NAT)

Network Address Translation can be seen as a bit of a 'workaround' for IPv4 to allow for a large number of devices to use a small number of public IP addresses available. As mentioned before, IPv4 allows for 4.3 billion addresses but there are over 7 billion people on the planet, each of which would probably have more than one device. To combat this, each device is given a private IP address. These addresses are private, no servers can communicate directly with a device using these addresses. The router acts as a proxy between the device and the router responds to the device.

For example, if you were getting ready to go to college in the morning and you see that the weather will be like later in the day, you could use your smartphone to look up the weather. The smartphone would send a packet to the router containing the private IP address (return address), the server (destination address) and a message of what is being requested. When the packet reaches the router, it creates a log in the NAT forwarding table and instead of broadcasting the private IP address, it will broadcast the network's public IP address. When the request is handled by the server, it formulates a response, the reply is sent back to the router IP address and the router then forwards the response to the device by checking the NAT forwarding table for the private IP address.



The outbound data request	The inbound data response
In the outgoing request, as described above, the smartphone sends a request message to the weather app server at 84.42.199.5 on a given port. In the request it also includes the return address for the data, but when the message reaches the router the return address is changed to the public address of the router.	If the router hadn't changed the return address, the weather app server would fail because the private address is non-routable. Instead, the router changes the return address to its public address and stores the private address in a table so that it can forward the response from the server to the device.

**COPYRIGHT
PROTECTED**




PORT FORWARDING

NAT sorts out the majority of problems when it comes to granting each computer on a network, but what if a computer wanted to connect to your network to use ports are used; they allow computers to exchange data directly between each other. This is a complex translation step that would increase the latency of the data connection.

You have already seen some well-known ports in one of the examples on previous pages. Port forwarding is similar to network address translation, but it operates on the port level. Once you have configured a router to allow you to put a port forward with a specific port number, so that all data requests for that port or are directed at that port will be forwarded to a specific computer on that network. This is often used for online games where a user has the opportunity to create a personal server for themselves. It is safer to enable port forwarding so that other users cannot connect to the server. This allows for complete access to the network.

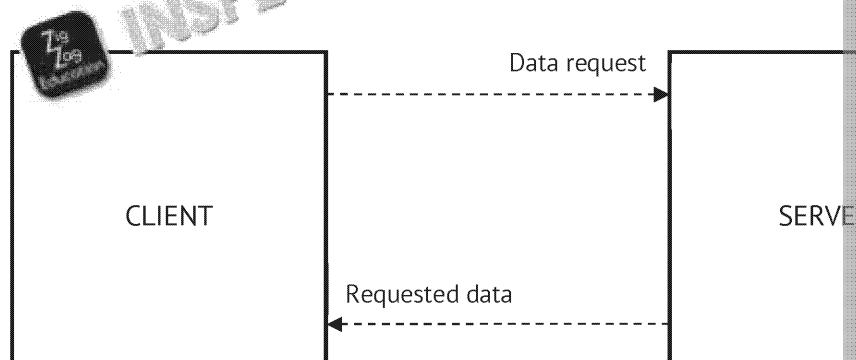
Take a look at the diagram below; the network has two internal connections and the router is a personal server for a game that a user has created and is on the network.



CLIENT-SERVER MODEL

The client-server model is a network architecture that dominates network design. In this model, computers in the network belong to one of two sets, or in some cases they can be both. A client is a machine or application that is usually interacting with the user. It is responsible for sending requests via the network to a server, or small group of servers. A server receives these requests, processes them accordingly and then returns the requested data to the client over the network.

The connection between the client and the server uses the 'WebSocket' Protocol, which is a programming interface over which a full duplex connection is established between the client and the server via TCP. This allows for a stable connection for data between both parties at any time.



**COPYRIGHT
PROTECTED**



The protocols used to control the data flow to and from the server use a protocol called Representational State Transfer (REST) which relies on the HTTP methods. The client cannot interact with the server through the REST API. Instead the server creates a set of instructions that take into account the functionality of any given method that needs to be available; these instructions form the REST API. Therefore, the REST API is created and run by the server. The client's web browser calls the API.

The acronym 'CRUD' describes the basic functionality that needs to be implemented by the REST API.

C – Create **R** – Read **U** – Update **D** – Delete

In order for the client computer to interface with a server database, the REST API methods to be mapped onto the basic SQL commands as follows:

HTTP Method	SQL Method
GET	SELECT
POST	INSERT
DELETE	DELETE
PUT	UPDATE

Any data that is transmitted during communication between a client and a server is either Extensible Mark-up Language (XML) or JavaScript Object Notation (JSON). A server admin might need to change some of the functionality; traditionally this is done using the XML language – there are some large differences between XML and JSON.

JSON is considered to be much easier to use than XML because:

- More compact (same complexity but operates much more easily)
- Easier to create during initialisation of the server
- It is higher-level – it is easier to read and understand
- Easier for the computer to perform the operations required, therefore more reliable

THIN- VERSUS THICK-CLIENT COMPUTING

In client-server situations a decision needs to be made as to who handles the data. This is a crucial decision to make because it can affect all clients connecting to the server. If you have a device or an early version of a smartphone then you can see the difference if you use the 'web version' of a social media site; the device can't handle the volume of data because the server it has connected to is intended for computers.

Thin-client computing

Thin-client computing is a new way in which a computer is used. 'Thin' computers have low processors and low amounts of memory that are continually connected to a server. This separates the computer into two parts where the thin computer performs the user interface containing only the minimum number of parts, whereas the server performs the data processing. If the server has any 'downtime' then all data transfer is halted, rendering the thin client unusable until the server reconnects.

**COPYRIGHT
PROTECTED**



A great example of thin-client computing can be found when looking at the Large Hadron Collider, the world's largest particle collider at CERN in Geneva. The collider consists of a 27-kilometre-long ring made of superconducting magnets that collides beams of particles together at 99% of the speed of light. There are so many sensors in the ring that it is impossible for a single computer to process the data that is created or to perform any operations on it, so data is stored and manipulated by a central server. Scientists can then request the results of the operations to be displayed on their monitors.



For more information see Sections 11 and Section 10.5

Thick-client computing

Perhaps unsurprisingly, the opposite of thin-client computing is known as thick-client computing. This is where the computer has the capability to operate and store data locally. This means that the data can be verified instantly, but this can also run the risk of data loss. This is more expensive than deploying a thin-client server but potentially expensive hardware to cope with the flow of data during connections that are intermittent than thin-client servers as there isn't the need for continuous data transfer.

One of the best examples of thick-client computing is that of online gaming. Modern games do not require a dedicated server; instead a single player will host the entire game but all of the calculations are done on the separate nodes (the client machines) which will connect to a central hub which will store various data such as the scores and locations of players. This is called for by other consoles. This can be shown when there is an element of lag or when there is an issue with the host's Internet connection. Everyone is affected and the game doesn't function properly; this is because the host can't communicate with the nodes and the game experiences 'lag'.

Questions: Transmission Control Protocol / Internet Protocol

- 1 Explain the following concepts:
 - a) IP address (2 marks)
 - b) Public and private addresses (2 marks)
 - c) Subnet mask (2 marks)
- 2 What is a MAC address? (2 marks)
- 3 What advantages does IPv6 have over IPv4? (3 marks)
- 4 What is Network Address Translation (NAT)? (3 marks)

**COPYRIGHT
PROTECTED**



10. Databases

Computer systems are nothing without data and information. The very purpose of a computer system is to produce the desired output. The processing of data is of the utmost importance in all areas of computer science. To learn about the database and how it can be used to become an efficient tool for storing and retrieving data.

This section covers:

10.1 Conceptual data models and ERM.....	p1	10.4 Structured Query Language.....	p3
10.2 Relational databases.....	p2	10.5 Client-server databases.....	p4
10.3 Database design and normalisation	p4		

10.1 CONCEPTUAL DATA MODELS AND ENTITY RELATIONSHIP MODELS


Logical data modelling is the process of using scenario information to produce a model for a database. Entities are abstractions of real-world concepts (e.g. Students) which are the data that needs to be kept in the system. Possible entities can be found by looking at the scenario. The relations between the entities can be found by looking for words that indicate how entities could form relationships and resolve any irregularities.

Databases are structured systems for holding data records. They are a step up from spreadsheets and tend to become less usable as the data structure becomes more complex. They are useful because of their usefulness as databases need to be easily searchable. They also need to be updated and maintained properly. Poorly designed databases can be very difficult to use so it is important to design them well. For this reason, conceptual data modelling is dependent on what other data.

Data structures can be modelled by breaking the proposed data structure down into its constituent relationships:

- An *entity* is any item in the system about which data is stored, e.g. *Student*.
- An *attribute* is a property of an entity, equivalent to a field in a database. For example, a *Manager* might be *Name* or *Age*.
- A *relationship* exists between entities. This dictates whether an entity has one or many relationships or just one.

A simple version of an ERD for the tutor example that follows would look like this:



INSPECTION COPY

COPYRIGHT
PROTECTED



Entity definitions

This is the act of producing an overview of the entities displaying all attributes and their relationships. It is a maintenance tool so that a database admin can see how the entities relate to each other. The attributes for a table are written in parentheses after the entity name. The first attribute is underlined) followed by all other attributes. The last attributes tend to be foreign keys and are marked by an asterisk (*). In the tutor example, the Tutor entity definition is:

Tutor (TutorID, FirstName, LastName, D_o_B, OfficeNo, Salary)

Questions: Conceptual Data Models and ERM

- 1 What are the 'top-down' models used for? (1 mark)
- 2 If the entity 'Salary' is defined to contain two fields, what might the entity definition be? (1 mark)

10.2 RELATIONAL DATABASES

This model states that all data is stored in relationships formed by tables. Data is stored in *tuples*, and the records are identifiable via their primary key. These relationships are capable of handling data while rejecting erroneous data types.

RELATIONSHIP TYPES

The relations in a database between entities are the foundations of the data model. A relation found between the entities where the variation comes from. Relations can be 'one-to-one' or 'many-to-many'.

One-to-one


Examples of non-redundant one-to-one relationships are relatively rare. The entities are linked to only one of another entity. Redundant one-to-one relations can be created by attributes, as this relation tends to only occur between attributes within the same entity. For example, if you look at the relationship between a tutor and their office,



One-to-many

One-to-many is the most common relationship type where one field from a table is linked to many fields from another table.

An example of a one-to-many relationship would be the relationship between a tutor and their qualifications. A tutor can have one or multiple qualifications, but a qualification doesn't make sense semantically or technically.




COPYRIGHT
PROTECTED



Many-to-many

Many-to-many relationships are quite common. For example, a tutor might have multiple publications. This type of relationship can become complicated because they cannot occur in real life; data cannot be transferred in that fast sense. Therefore, many-to-many relationships are resolved by using an *association*. For example, a tutor can publish many papers and a paper can have many publishers so the association is many-to-many. This is because it's the simplest link.



ATTRIBUTES AND KEYS

Attributes

These are the elements of data that denote properties of the parent entity – the *attributes*. For example, a table that needs to hold personal information about the tutors in a database would include name and contact details but you'd need to consider items such as office number. Look at whether the attributes you store can be derived by other attributes, for example, of birth or age, not both.

Primary keys

A primary key is an attribute that allows others to be uniquely identified. For example, a primary key in our example would be first name. However, a primary key needs to be unique. Instead, we create a new field, one that allows us to create complete identification (e.g. TutorID).

Composite key

A primary key that involves more than one attribute is known as a *composite key*. For example, a composite key would be one where the names of people are stored as first name and surname. First name or a surname alone would not be sufficient as a primary key. For example, 'Mark Jones', 'Mark Peterson' and 'Jan Peterson' which means that neither the first name or surname would identify a person individually. In order to identify them uniquely you need both names.

Foreign key

A foreign key is a primary key in another table. Flat databases don't use foreign keys. In reliable and robust databases relationships *must* be created. These relationships prevent duplicate data from the tables by referencing to existing items. In the tutor example, if a tutor forms a department, the Department Code will become a foreign key in the tutor table.

Questions: Relational Databases

- 1 What is meant by a relational database? (1 mark)
- 2 The following definition is for an employee entity within a relational database.

Employee (EmployeeID, FirstName, LastName, D_o_B, OfficeNo)

 - a) Identify the primary key in the Employee entity. What is a primary key?
 - b) Identify the foreign key in the Employee entity. What is a foreign key?

**COPYRIGHT
PROTECTED**



10.3 DATABASE DESIGN AND NORMALISATION TECHNIQUES

Normalisation is a rational data analysis tool useful for validating entity relationships and optimising database structure. In order to normalise, several (as many as possible) forms of the form of data samples. These are inputs (forms) and outputs (reports and queries).

UNF – Un-normalised Form

1. Select an *initial key*. This acts as a ‘starter’ key. It needs to be unique and you can derive other values from it. If there is no suitable key, add one.
2. Transfer all identifiable attributes, ensuring each has their own *relevant* and *unique* name.
3. Look for repeating groups. These are a group of similar attributes that have multiple values for a single value of the initial key. Select a suitable attribute for the repeating group, surround them with brackets, and write them down from the initial valued attributes.

As you can see in the example on the right, the group in the brackets is the repeating group and is separated from the non-repeating group. SkillCode is promoted to become a *key attribute* for the relation. Note that in un-normalised form there is no propagation for the separated group.

1NF – First Normal Form

This is the most basic of the normalised forms; in order for data to be in the first normal form you must:

4. Create new relations by separating all repeating groups; select a new initial key for the new relation and *propagate* (copy) the initial key to form a *composite key*.
5. All other single-valued attributes remain with the initial key.

As you can see below, this is the first step to providing some structure and relationships to the database. There is a new composite key formed when you propagate the initial key down to form the new relation with SkillCode.

2NF – Second Normal Form

The second stage of normalisation requires that all key attributes in a table are independent of each other and that the table is in first normal form. Only then can you move on to creating a database that is in second normal form which will provide more strength to the data structure.

In order for the data to be in second normal form you must:

6. Separate attributes from keys formed in the previous step that are *transitively* dependent on one part of the composite key.

As you can see in our example to the right for the tutor example, SkillDescription is only dependent on the SkillCode, not the TutorID. If this was left in first normal form and implemented as a solution, you would quickly find that the database was not performing as optimally as it was intended to and would accumulate redundant data.

This would slow the database’s performance greatly. Logically it makes sense to create a new relation to prevent this from happening. Key propagation must be done again but only for the original primary key for that relation (SkillCode). There will have created a new relation as can be seen on the right.

**COPYRIGHT
PROTECTED**



3NF – Third Normal Form

The third and final step of the standard normalisation is the transition into the third normal form. This is the step that most people struggle on, but if you use the pointers and look at the example that has been used throughout the steps it is clear to see what is required of you. In order for data to be in the third normal form you must:

7. Separate any attributes that are dependent on other non-key attributes. If any foreign keys are retained in the original relation.
8. Check composite keys for redundant parts. If a part of a key can be determined from other attributes, demote key attribute to non-key.

Take a look at the example on the right: in the first group that had been left unchanged since the beginning we have moved two attributes down and created a new relation for each. This has been done because for each they are only partly dependent on the Tutor ID attribute. If there would be multiples for each. There will be multiple tutors on the same department faculty and there will be multiple tutors on the same pay grade. You will also need to create a new 'ID' attribute for each of the new relations and use these as foreign keys in the first relation. The database is now in third normal form.

For an overview of all the steps take a look at the table on the following page. If you work from left to right on the table you can see how each step naturally and logically leads to the next, and how the final format of the data is optimised for removing redundant and duplicate data. Following this table format can be very beneficial as the information you need is available for you and you can see how entities begin to

Did you know?!

The normalisation ladder actually has six more forms before the database is fully normalised. These push the database theories as far as the theory is concerned, but the final steps are aimed towards improving the performance of joins together during a query or maintenance of the database. Seeing as joins can be very slow on the size of the database, these final steps are seen as the final requirement but only for applications where the database needs to operate as smoothly as possible.

You'll be pleased to know that this is well and truly out of the scope of this course!

**COPYRIGHT
PROTECTED**



Normalisation example – overview

Salary and department are separated and their ID fields are retained as foreign keys.
The entity names have been added to the right.

UNF	1NF	2NF	3NF
<u>TutorID</u>	<u>TutorID</u>	<u>TutorID</u>	<u>TutorID</u>
FirstName	FirstName	FirstName	FirstName
LastName	LastName	LastName	LastName
D_o_B	D_o_B	D_o_B	D_o_B
Grade	Grade	Grade	OfficeNo
SalaryCode	SalaryCode	SalaryCode	SalaryCode*
OfficeNo	OfficeNo	OfficeNo	DepartmentCode
DepartmentName	DepartmentName	DepartmentName	<u>DepartmentCode</u>
(SkillCode	<u>TutorID</u>	<u>TutorID</u>	DepartmentName
SkillDescription	<u>SkillCode</u>	SkillCode*	<u>SalaryCode</u>
Qualification)	SkillDescription	Qualification	Grade
	Qualification	<u>SkillCode</u>	<u>TutorID</u>
		SkillDescription	SkillCode*
			Qualification
			<u>SkillCode</u>
			SkillDescription

Bold & underline indicates primary key

Asterisk (*) indicates foreign key

Normalisation is primarily used by system and data analysts. When carried out correctly, it results in an optimised database model. The resulting structure means that managing the data is easier. The span of the solution is increased. The relationships will eliminate redundant data, making changing data easier. In the tutor example, if a tutor changes address or department, only one record would need to be changed, but if the database is not normalised the *database model* would need to be changed for all changes.

Questions: Database Design and Normalisation

- Describe the steps of transitioning from the following:
 - 1NF to 2NF? (1 mark)
 - 2NF to 3NF (2 marks)
- How does normalisation affect a database's performance and reliability?

**COPYRIGHT
PROTECTED**



10.4 STRUCTURED QUERY LANGUAGE (SQL)

Structured query languages are the regular languages that are used to create databases in a variety of management systems. There is more than one type of 'SQL' but the original language that was developed in the early 1970s is called 'SQL'. The functionality still remains the same. The aim of SQL was to provide a language that was very high level so that database administrators could manage

THE DATA DEFINITION LANGUAGE AND DATA MANIPULATION LANGUAGE

There are two language sections in SQL that have been developed for performing database operations; these are the *data definition language (DDL)* and the *data manipulation language (DML)*.

DDL

The *data definition language* is a language used to build the structure of a database. It consists of commands which allow the user to define the structure of the database by creating tables, indexes, and views. The DDL also allows the administrator to apply constraints to a database table, such as the constraint of access rights for a table to prevent certain users from editing data altogether. Here is an example of a DDL statement to create a database called 'customer' and 'order':

```
CREATE DATABASE db; # create an empty database
CREATE USER dbuser IDENTIFIED BY 'password123'; # identify user
GRANT ALL ON db.* to dbuser; # grant all permissions to user
CREATE TABLE db.customer ( # create a new table in db called 'customer'
    customerID INTEGER, # attribute definitions (name and data type)
    firstName VARCHAR(20),
    surname VARCHAR(20),
    phone VARCHAR(14),
    PRIMARY KEY (customerID),
    UNIQUE INDEX (customerID)
);

CREATE TABLE db.order ( # create a new table in db called 'order'
    customerID INTEGER, # attribute definitions (name and data type)
    orderID INTEGER,
    PRIMARY KEY (orderID),
    FOREIGN KEY (customerID) REFERENCES db.customer(customerID),
    UNIQUE INDEX (orderID)
);
```

Note:

The 'CREATE USER ...' and 'GRANT ...' commands simply create a user with a password and allow them to do whatever they want with the database. The 'CREATE TABLE <table_name> (<column1> <data type>, PRIMARY KEY (<columnx>));' construct is probably the most common one as it is how most databases are defined using this language. Another significant command is 'DROP DATABASE' and 'DROP TABLE' to delete databases and/or tables.

The normalised database with multiple tables is linked together utilising the 'FOREIGN KEY' command shown above. This links the tables by identifying that the 'customerID' in the table 'customer' is the same as the 'customerID' in the table 'order'.

INSPECTION COPY

COPYRIGHT
PROTECTED



DML

The *data manipulation language* is the language used to populate and update tables created. DML allows the user to insert, modify, delete and query a database table, for a specific record or for records that meet criteria. Here is an example

```
USE db;
INSERT INTO customer(customerID, firstName, surname, phone)
VALUES (1, 'Bob', 'Smith', '+441234568909');
```

The basic commands for the DML are:

SELECT	Retrieves the data of a record from a table
INSERT	Inserts data into a table into a new record
UPDATE	Overwrites existing data in a record
DELETE	Removes a record from table

It is these commands that will allow you to implement a database, perform operations on data and retrieve information from a database.

QUERYING DATABASES USING STRUCTURED QUERY LANGUAGE

Structured Query Language (SQL) is a language designed to search databases for information.

In order to gain some understanding of SQL we will look at some of the queries we will call Customers.

Surname	Age	Children	Occupation
Harvey	21	1	Salesman
Ellison	67	2	Retired
Watson	40	6	Nurse
Clarkson	25	0	Unemployed
Butcher	18	0	Butcher
Smith	30	4	Doctor

Similar to programming in an imperative language, SQL uses comparison operators.

=	Equal to	>	More than
<> or !=	Not equal to	<=	Less than or equal to
<	Less than	>=	More than or equal to

**COPYRIGHT
PROTECTED**



LIKE

The LIKE operator allows you to search for wild cards that come close to matching the query. Take a look at the

LIKE	Result
LIKE "D*"	This will search for strings that begin with the letter 'D' followed by any number of characters.
LIKE "*son"	This will search for strings that end in the string 'son' preceded by any number of characters.
LIKE "*l*"	This will search for strings that contain the letter 'l' preceded or followed by any number of other characters.
LIKE "Sm?th"	This will search for strings that match the search string 'Sm?th' where '?' represents any single character.
LIKE "#3"	This will search for a string that varies by a character followed by a numerical value as opposed to a string value.

Given what you've been told about the DML above, you will now see a simple example of how to use the SELECT command. The SELECT command selects which columns to return, the FROM keyword selects which table the results should be drawn from. The WHERE keyword is optional and basically means 'all'. The statement below shows how this would be used.

```
SELECT Surname, Occupation FROM Customers
```

This statement would return:

Surname	Occupation
Harvey	Salesman
Ellison	Retired
Watson	Nurse
Clarkson	Unemployed
Butcher	Butcher
Smith	Doctor

WHERE

However, it is very rare that you will require all the information contained in a table. You can accomplish this by using the WHERE keyword. For example, you could retrieve the records of customers that have no children.

```
SELECT * FROM Customers WHERE Customers.Children = 0
```

This statement would return the following:

Surname	Age	Children	Occupation
Clarkson	25	0	Unemployed
Butcher	18	0	Butcher

**COPYRIGHT
PROTECTED**



Conditionals can also be used. For example, to select everyone 40 or more years old, the following statement would be used:

```
SELECT * FROM Customers WHERE Customers.Age >= 40
```

Other keywords also exist which allow you to select a particular range of values. For example, the following statement would select everyone who is not married and does not have any children:

```
SELECT * FROM Customers WHERE Customers.Age > 20 AND Customers.Married = 0 AND Customers.Children = 0
```

Another useful clause to use is ORDER BY. This means that the results are returned in either ascending (using the keyword ASC) or descending order (using the keyword DESC). Take the following as an example:

```
SELECT Surname, Age FROM Customers WHERE Customers.Children = 0  
ORDER BY Customers.Age ASC
```

This would return:

Surname	Age
Harvey	21
Smith	30
Watson	40
Ellison	67

Multiple tables in SQL

As has been seen with normalisation (section 10.3), the data in a database is often spread across multiple tables. It is important to be able to retrieve, update and delete data using other references.

For example, using the orders system it may be important to contact the person who has placed an order to tell them it has arrived or similar. The order ID is in the orders table, whereas the customer's name is in the customer table.

Method 1

```
SELECT customer.Surname, customer.FirstName, customer.phone FROM customer, orders  
WHERE orders.orderReq AND customer.customerID = order.CustomerID
```

Method 2

```
SELECT customer.Surname, customer.FirstName, customer.phone FROM customer, orders  
WHERE orders.orderReq INNER JOIN customer ON customer.customerID = order.CustomerID
```

Please note that the SQL here shows both tables in the FROM statement; some versions of SQL only require just the 'customers' table with the inner join referencing the 'orders' table.


You will be required to be able to write and interpret SQL statements in an examination. You should practise these statements. There are several ways of doing this: you could set up a small database on some network space or local host; some DBMS software such as Microsoft Access can be used to create query definitions.

**COPYRIGHT
PROTECTED**



10.5 CLIENT–SERVER DATABASES

The following diagram is the same as the one found in the notes for *client–server* databases. It shows the basic structure and actions of the model, and the same applies for a client–server database.



The client can be a node on a network, an individual computer or someone's laptop on a planet and the structure still works the same. The idea is that a client requests data from a server by a connection made over a wide area network (such as the Internet) or a local area network (such as those found in an office building). Once the data request is fulfilled, the connection is terminated. An example of this is found on any Internet browser; the browser acts as a client computer with a database that has the stored data of the website, and once it receives the data packets to display a website the connection is terminated.

Another example can be seen in online banking. When online banking was first introduced, there were many questions and concerns to be answered about security, integrity and cost. These concerns were most often unfounded as most do not understand how the technology works. The client acts as the interfacing client for the database again and sends a request for data from the server. The server replies with the authentication and authorisation steps that are required to request the information. If the checks are passed then the request for the data is fulfilled and the server returns the data that was requested, and the connection for that specific request is terminated.

ADVANTAGES AND DISADVANTAGES

Advantages

- Centralisation (data being stored in a single location) – This allows server administrators to manage the server's network, how the data requests are handled and how files are stored.
- Accessibility – As long as the server is up and running then the data is available at any time and from any location as long as the user has permission to access it.
- Adaptability – The model allows for the use and adaptation of the technology to suit the needs of the user (for information see Section 10.6).
- Scalability – Due to the adaptability of the model, the database can be scaled to meet the needs of the user and it doesn't necessarily mean that the client's computer power is increased; only the database's hardware is increased.
- Backup and recovery – Because all the data is centralised it means that backups of the database can be made so that if there is a failure or loss of data the data can be recovered with no consequences for the users of the service, minus server downtime.
- Security – The model allows a server/database administrator to control access to the data and how data is returned to the client. Limiting what the client can do with the data in the database can be preserved. This is also increased when the connection between the client and the server is dropped after the request is fulfilled.

INSPECTION COPY

COPYRIGHT
PROTECTED



Disadvantages

- Network overload – Due to the way the database is implemented, it can become overloaded with data resulting in slow response times or *concurrent access*. This can lead to data loss, corruption or *concurrent access*.
- Cost – These servers and their hardware are not cheap to implement and the cost is increased when the server is being used over the long term or for frequent updates. There is a need to prevent a network overload and server downtime.
- Robustness – It is the centralisation of the data in the database that makes it less robust as a peer-to-peer network. In peer-to-peer networks, if the server goes down, the performance of the data can be transferred to another host computer. In a client-server model, if the server in a client-server model is lost then all transitions of data terminate until the server comes back online, which may require special maintenance.
- Maintenance – There is a level of knowledge that is required to maintain a server. This need is increased when you're dealing with a service that many people are using. This is why the database isn't available (downtime). A specialist technician, is required to ensure that the network isn't becoming overloaded and to perform maintenance tasks on the server's hardware.

Although the client-server architecture is costly to maintain, it is relatively easy to use and is ingrained into modern life and how we interact with data because of how it allows access to the server. *Concurrent access* was referred to above; this is when the server allows multiple users to access the same piece of data at the same time. However, *concurrent access* can occur when multiple users access the same piece of data at the same time, meaning that an update is lost between users.

CONCURRENT ACCESS

Concurrent access is making sure that more than one user can at least view the same data at the same time, and there can be many reasons as to why you would want this in a computer system.

However, the real issue with concurrent access is that it could lead to both users making changes to the file, in which case one change will always be lost after both users save the file. This will inevitably lead to data loss from a system and potential critical errors for the clients connecting to the server.

Protocols have been developed to help maintain the integrity of the data in the database if concurrent access occurs. Concurrent access can be prevented by using a lock.

Record locking

Record locking is something that your computer will implement when you have a file open and another copy of the file. The second instance of the file will be opened in a read-only mode, meaning that no changes to be made. This is also implemented on a database and is the act of locking a record so that it cannot be accessed while someone is editing the file or data item.

Serialisation

Another method of preventing concurrent access is to create a clone of a data item, allow the user to apply changes to and then upload a copy of this clone to the database. This is called *serialisation*, and by creating a clone of an object before amending changes it can be lost during editing.

Timestamp ordering

Timestamping is a non-lock method of concurrent access so that multiple people can access the same data at the same time without causing ordering errors. The assumption is that all transactions have a unique timestamp so multiple users can access the database data.

**COPYRIGHT
PROTECTED**



11. Big Data

Big data is seen as an 'all-encompassing' term given to data that won't fit the usual data constructs or containers.

The reference to 'big' however is potentially misleading as it implies that volume is the only factor when determining whether data should be classed as 'big data', when in fact there are factors to consider.

These factors, and some ways of addressing the issues posed by Big Data are explored in this section.



Big data can be described in terms of:

- **Volume** is the obvious reason why data would be considered 'big'. If it is no longer suitable to store on a single server then there is no process that data without breaking it down somehow. You may think it would be to store it across multiple servers on a relational database. However, data does not scale very well across multiple machines.
- **Variety** is possibly one of the more relevant reasons to call something 'big'. Data comes in different formats and data types that it becomes difficult to store and hard to interpret meaningful information.
- **Velocity** is the speed at which data needs to be accessed, i.e. during a short time to respond to the data.

Did you know?!

If you fully compressed all of the new data created daily, how much informational data storage do you think is required? A few hundred gigabytes? Terabytes?

Not even close, as of 2014, 2.5 exabytes of information is generated each day. That's 2.5 billion gigabytes, 2.5×10^{18} bytes, every single day.

The largest data capacity facility is found at the Utah Data Centre (shown here on the right) and can store a massive 12 exabytes of data.



VOLUME, VARIETY / VARIATION, VELOCITY

Volume

Volume of data are the easiest to picture when it comes to handling data because everyone can relate to how difficult it is to handle and cope with large amounts of information. When a dataset is so large that it can no longer be accommodated by a single database then it is easy to see why it is given the name *big data*. It is hard to link computers together in such a way that the memory in the computer can function as a single unit, so instead the dataset is broken down and spread across large arrays of computers in a database and the data can then be processed.

**COPYRIGHT
PROTECTED**



To help put big data into context, consider the following. At CERN in Geneva contains 150 million sensors, each with a polling rate of 40 million times every second of *raw* input data created every *second*; in a single day it would fill 86 1 terabyte calculation, that is 30 petabytes (31,457,280 GB) of data being recorded annually. All of that raw data needs to be stored before it can be processed and that is where big data comes into use; this is explained *Section 13*.

Variety

There is no way for a computer to visualise a dataset and see patterns without being able to look at each individual element of the dataset. The odds of a dataset being retrieved in a logical order that is ready to be processed without being stored and processed first are minimal, so it is natural to assume that there is always going to be variety in a dataset especially when it is raw data that is being collected.

Within a dataset there can be large differences in the data that is being processed. This is why it is one of the more relevant constraints of big data – that a system can struggle to cope with large quantities of data that varies in such a way that it becomes difficult to know how to store it. On large networks where there is the need for varying types, it is easy to see how a computer system can be weighed down

Velocity

Velocity takes into account the increasing rate at which data is transmitted through a network, and it has followed the same trend as the volume of data that is created each year.

In the past, our networking capabilities were quite primitive compared to the fast networks that we can develop today. These primitive networks restricted the way in which we could access data by restricting the speed at which we could access it. There have always been situations where a fast computer was needed to turn out and move data quickly and efficiently.

For example, the stock market servers must be ultra-efficient because of how frequently the stock prices and resources change.

It's not just where fast computing is needed, though; the way in which the globe has changed has had a large impact on how data is handled. When you load a page on your smartphone, the data has to be collected from a server on the network provider's mast. Imagine how many people are connecting to the mast at any one time – that is required – but many aren't satisfied with slow data connections and so new technology has had to be developed to enable this.

**COPYRIGHT
PROTECTED**



ADDRESSING THE ISSUES

When data banks become so large that they will no longer fit onto a single server, data must be distributed across a bank of computers. Using more than a single server for a data bank requires specialist programming that is very complex and expensive to make to order for the data bank that is being processed. However, the functional programming approach (*Section 12*) can be used because it makes it easier to create and maintain code that can be efficiently distributed, because it supports:

- *Immutable data structures* are data structures that are unchanging. The data that is used cannot be changed or altered, meaning that there are strict rules about how data can be manipulated.
- *Statelessness* is inherent in functional programming – that is, the part of a program that is not of state; it doesn't remember the results or states of any preceding execution. This means that an instruction being executed, and states can sometimes be restricting. For example, a thread enters an altered state because a core has multiple threads running. This means the program could run over many threads.
- *High-order functions* are functions that can accept a function(s) as an argument and return a result. This allows the language to be highly adaptable for whatever purpose it is used for.

These features allow a programmer to produce the code to handle vast volumes of data, which means that the code will be correct for its purpose.

Fact-based model

Generating a fact-based system is much like creating the Unified Modelling Language (UML). If you've not come across UML before, it is a set of approaches for representing the functionality of a system in a way that can convey large systems in a relatively comprehensive manner. Similarly, fact-based systems are not concerned with how the dataset is structured, how the dataset is linked and how the dataset can be used. They are concerned with how to handle vast quantities of data without needing to be concerned with how the data is structured. The system will only perform operations within the constraints of the facts that it is given.

The databases that you're most likely to have experience with aren't very efficient at handling large quantities of data due to how they are constructed. There is an efficient way to handle large quantities of data and that is by using what is known as a graph database.

Graph databases

Graph databases offer the same functionality as a standard database (see *Section 11*). Instead of being an index for a field or an entity, the graph database uses a pointer to link data. Graphs are an abstraction of data that are related to each other or linked and are used to represent *relationships* and *properties*; these three criteria are found in the database's *schema*. A schema is a description of how the data is structured, how the data is stored and what constraints there are on the data. It is a way in which data can be stored and retrieved.

Nodes

In a relational database the data of an instance of the attributes is stored in a table. Each record contains all the data values for the attributes in the data table and it is these records and the data tables are linked to each other to form relationships. In a graph database the elements of data themselves are stored in what is known as a *node*. Each node is a representation that describe the object in its entirety.

**COPYRIGHT
PROTECTED**

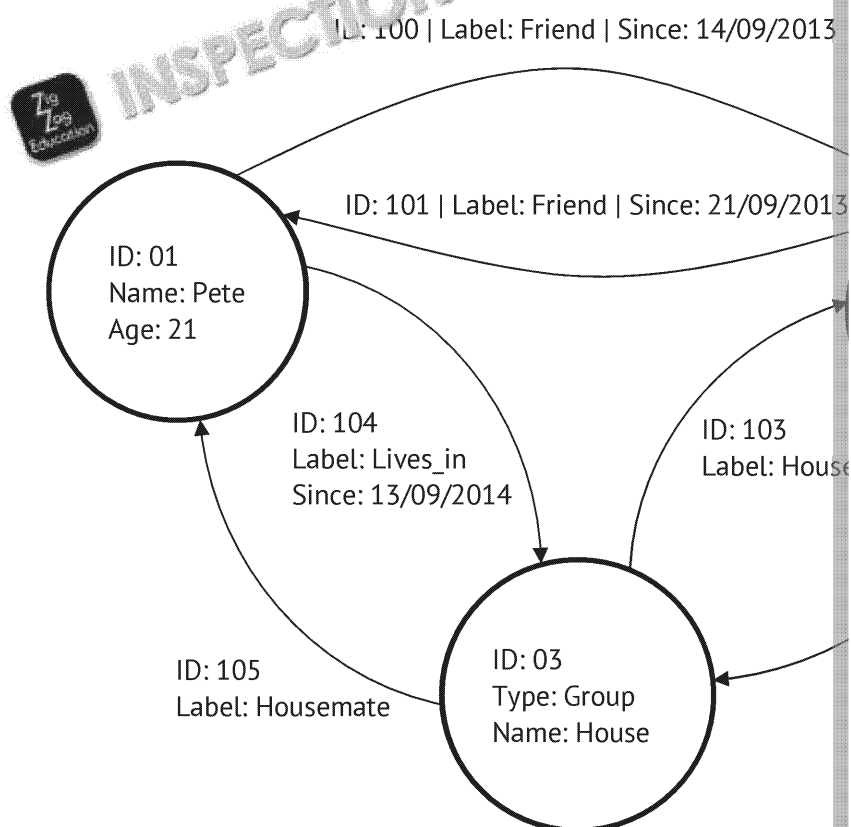


Properties

In a relational database the properties of a record are assigned in the DDL and contribute to the database's structure. In graph databases the properties of a *node* or the *edges* between nodes and are represented by text. The properties are attached to, including the name, attributes and how the items are linked.

Edges

In a relational database the relationships between tables are shown in the *relationships* table, how the links between the tables will provide integrity and how the data stored is represented by the *arrows* between graph objects (nodes). In graph databases the relationships and instead the relationships between nodes have properties of their own and



Questions: Big Data

- 1 What is 'Big Data' and what three criteria can define it? (3 marks)
- 2 How does functional programming help to address issues posed by Big Data?

INSPECTION COPY

COPYRIGHT
PROTECTED



12. Fundamentals of Functional Programming

Functional programming is called such because its primary and fundamental approach is to use functions. In this section demonstrates how functional programming can be used to solve problems.

This section covers:

- | | |
|--|--|
| 12.1 Functional programming paradigmp1 | 12.3 Lists in functional programmingp5 |
| 12.2 Writing functional programsp3 | |

12.1 FUNCTIONAL PROGRAMMING PARADIGM

Just as with other paradigms, you can create a function body that you can pass parameters to. In functional programming the entire body of the program is a function that takes arguments and returns a result which is returned. Generally speaking, the paradigm is one that encourages writing programs using complete computation by using functions as arguments for the function and returning a result, rather than the explicit coding found in paradigms such as imperative or object-oriented.

There are programming languages you can learn that are only applicable to functional programming, but you can just as easily use the programming language you've already learnt if it allows functional programming. Most modern multi-paradigm languages do. You will be expected to show the ability to write simple programs in the paradigm to perform a simple task, which could be as simple as calculating the average of a set of numbers in an array to writing a function that produces the average of a set of numbers. The code in this section will be written in *Haskell*, a *purely functional* programming language.

FUNCTION TYPE

Functions in all languages must have a type, often called their return type. What this means is how the function is mapping the input and creating a return. It is said that a function f has a domain A and a co-domain B .

$$f: A \rightarrow B$$

This says that the function f , has a domain A , and a co-domain B . This shows the input argument for a function and the co-domain is your returned type. When the input type is A and the result type is B . There are some limitations, though; in order for a function to be a function, A and B have to be subsets of objects of some data type.

In Haskell, a function's argument and return type must be declared.

For example, if you wanted to write a function that added two numbers and returned the result, the function `addNumbers`:

Haskell

```
addNumbers :: Double -> Double -> Double
addNumbers x y = x + y
```

This may seem a little confusing, but we define the output as well. So in this case, the function `addNumbers` returns a double!

FIRST-CLASS OBJECTS

First-class objects are actually a concept that occurs in many programming languages. You've probably already witnessed without noticing. First-class objects are objects that can be created in an expression, be assigned as an argument, be assigned as a variable or be returned from a function. In most programming languages first-class objects are integers, floating point, Boolean, and strings.

A first class function can return and take in functions. You will see an example of this in the next section.

**COPYRIGHT
PROTECTED**



FUNCTION APPLICATION

Unlike in traditional coding paradigms that you're probably used to, the way in programming is slightly different. In functional programming, to apply a function, you apply it to its arguments to produce the result. The process of giving specific values to a function is called *application*. This is how the entire paradigm works; it is the application of functions to values. A function will return a value. For example, if we had written *our addNumbers* function, but

Haskell

```
addNumbers :: Integer -> Integer -> Integer
```

The function will continue to work as long as the input types of the input match. It will work for decimals any more as the function has been built to handle decimal values. The function argument as a pair of values: the function cannot work or be called without two arguments. *addNumbers* (4) would not work. To add 4 + 4 we would call the function *addNumbers* 4 4. This function is a pair.

PARTIAL FUNCTION APPLICATION

In functional programming, literally everything is a function, and like all functions, they return a value. However, with the functional paradigm all functions actually return a function. While this sounds as comprehensive as the sound of one hand clapping it is a key concept to handling data. What the function actually does is return a function that actually does the work, called *currying*. Take a look at the following example.

The function *myMult* takes two arguments: *x* and *y*.

```
myMult x y = x * y
```

The calling scheme of the function could be written as:

```
myMult :: Int -> Int -> Int
```

Partial functions can also be defined as functions whose result is only part of a problem. Problems can occur and needs to be investigated. For example, the function *half* of the domains of natural numbers (\mathbb{N}) and real (\mathbb{R}).

However, the function *half* $x = x/2$ is defined with the domain of real numbers.

```
half :: real -> real
```

However, it does not work with natural (\mathbb{N}) if the number is odd. So the domain of natural numbers would have to be defined using the domain ($2\mathbb{N}$), i.e. the group of even numbers.

COMPOSITION OF FUNCTIONS

Functional composition is the process of combining two functions to produce a new function. It is an aggregation of both in one function, which appears to be a single task. This is something that is common in mathematics and is something that you're taught to do from the start. It is a natural progression that you'll undergo as you learn the ins and outs of programming. In C# you can declare an array and then populate it, or you can do both steps in one line, which is more succinct. It is similar in functional programming to look for a quick way to do things.

Suppose you wanted to remove the top element from a list, and then reverse the list. In functional programming, functions exist as functions, so using function composition we can do it in one go:

Haskell

```
reverse . tail [1,2,3,4,5]  
output: [5,4,3,2]
```

Here you can see the dot notation combining functions together, and the result is *reversedSorted* which is then printed.

**COPYRIGHT
PROTECTED**



12.2 WRITING FUNCTIONAL PROGRAMS

FUNCTIONAL LANGUAGE PROGRAMS

Before you get too concerned, the level of your programming skills for functional programming is to be near the level of your imperative programming skills. In this section you will build on your knowledge of the functional programming paradigm and it will aim to give you a solid foundation to then build on. Make sure you read the code slowly and read the comments carefully.

There are programming languages that were developed solely for the purpose of functional programming. It is also a feature that is being introduced to many modern multi-paradigm languages.

As mentioned before, one of the most important features of functional programming is the concept of a 'function'. A function is *high-order* if it takes a function as an argument or returns a function. The three that you're expected to know are *map*, *filter* and *reduce (fold)*.

Map

In its simplest form *map* is a function that accepts a list, *L*, as an input whose elements are of type *A*, and a function, *f*, which maps *A* to another type *B*. *map* then applies *f* to each element of the list and returns a list of the results as type *B*. A very simple Haskell example can be seen below.

Haskell

```
square :: Integer -> Integer
square x = x^2
map square [1,3,5,7,9]
```

As you can see, first you define what the function, *square*, is doing and then you use it to produce more advanced programs. The list would be passed in as a parameter to the *map* function. The return of the function would be *1, 9, 25, 49, 81* as *map* applies the *square* function to each element of the list.

Filter

Filter uses a Boolean value called a *predicate*, and a list as inputs. The function applies the predicate to each element of the list and returns a list of all elements that satisfy the predicate. The predicate is a function that takes an element of the list and returns a Boolean value. The predicate is used to decide if the element should be added to the newly created list.

Haskell

```
filter odd [1,2,3,4,5]
```

In Haskell, 'odd' is a function that returns true if a value is odd.

Reduce

Reduce, also known as fold, is a function that is a way of reducing an entire list to a single value, *B*, through built-in functions. It returns a single value as a list. A simple example is to reduce a list of numbers and produce the sum of the list. In Haskell there are two functions, *foldl* and *foldr*, which act slightly differently. *Foldr* works by evaluating the rightmost element first, then moving left. *Foldl* works more using iteration, then building up the result in a similar fashion to recursion. *Foldl* works more using iteration, then building up the result in a similar fashion to recursion. *Foldl* works more using iteration, then building up the result in a similar fashion to recursion.

Haskell

```
foldr (+) 1 [2,3,4,5]
-- (1+(2+(3+(4+5))))

foldl (+) 1 [2,3,4,5]
-- (((1+2)+3)+4)+5
```

COPYRIGHT
PROTECTED



12.3 LISTS IN FUNCTIONAL PROGRAMMING

LIST PROCESSING

When processing lists in a functional programming language, you can talk in terms of a list written as Head: Tail. For example, the list [10, 9, 8, 7] has a head of '10' and a tail written as 10: [9, 8, 7]. Similarly to how list structures are written in other languages, an empty list is completely empty and is written as '[]'.

Functions of a list

Here are the functions which you need to know when working with lists. You can think of lists as instead of lists, as really they are just lists of linked factors!

Returning head of a list

This is similar to searching for a given index in a list. 'Head' is a keyword in terms of a list, and is shown in the following:

Haskell
head [1,2,3,4,5] output: 1
head "Hello" output: 'H'

Returning the tail of a list

Almost identical to returning the head value of a list, this function returns the tail of a list, as shown in the following:

Haskell
tail [1,2,3,4,5] output: [2,3,4,5]
tail "Hello" output: ello

Test for an empty list

This is a very basic test that can be used to prevent underflow errors in a program.

Haskell
null [1,2,3,4,5] output: False
null "Hello World!" output: False

Return the length of a list

The length of a list can be returned using the *length* function as shown:

Haskell
length [1,2,3,4,5] output: 5
length "Hello World!" output: 12

Note the answer to length "Hello World!" as it includes the space between the two words.

**COPYRIGHT
PROTECTED**



Prepend an item to a list

In order to prepend (adding something to the beginning of a list) you must create a new list, then assign it to the list using colon notation. Take a look at the following example:

Haskell

```
li = [1,2,3,4]
st = "Hello world!"
```

```
3:li
output: [3,1,2,3,4]
```

```
"I want to say" ++ st
output: I want to sayHello world!"
```

Note that as the space was missing in the first list, the append adds directly to the end of the list (see below) rather than prepend. For strings it is a little bit different as it can only add one character at a time.

Append an item to a list

To append an item to the end of a list you use exactly the same principle as we have used to prepend. To append an item to the end of a list you use exactly the same principle as that we have used to prepend the same variable to hold the value of the new element.

Haskell

```
st = "Hello world"
st ++ " out there!"
output: Hello world out there!
```

Combining lists

Just like text, we can add lists by using the ++ function so:

```
[1,2,3] ++ [4,5,6]
```

Combines to

```
[1,2,3,4,5,6]
```

Finding elements of a list

We can find an element of list using !!

```
[1,2,3,4,5] !! 3
```

This will find the number 4. NB remember that list elements count from 0 so

```
[1,2,3,4,5] !! 0
output: 1
```

This also works for text:

Haskell

```
"Hello world!" !! 7
Output: 'w'
```

```
"Hello world!" !! 0
output: H
```

**COPYRIGHT
PROTECTED**



Mathematical functions on lists

You can build your own functions to work on lists, but basic calculations are:

Add up a list	Haskell
	sum ourList Output: 15 sum ourStrList Output: Error

Multiply a list	Haskell
	product ourList Output: 120 product ourStrList Output: Error

Arrange elements to largest	Haskell
	sort ourList Output: [1,2,3,4,5] sort ourStrList Output: " HWdellloor"

Note the space is
capitals are before

Reverse the order	Haskell
	reverse ourList Output: [5,4,3,2,1] reverse ourStrList Output: "dlrow olleH"

Find minimum value	Haskell
	minimum ourList Output: 1 minimum ourStrList Output: ""

Find maximum value	Haskell
	maximum ourList Output: 5 maximum ourStrList Output: ""

Custom functions on lists involve function declarations.

Double every element	Haskell
	[x*2 x <- ourList] Output: [2,4,6,4,10]

Square every element	Haskell
	[x*x x <- ourList] Output: [1,4,9,16,25]

By using the above functions, it can be seen that by combining and working with lists, many different results can be produced.

**COPYRIGHT
PROTECTED**



Recursion using lists

We can use the head and tail of a list to write functions that use recursion.

Suppose you wanted to double every element of a list without using the map function. You could do it as above, or you could write a function that goes through each element and doubles it.

In Haskell we can do this using the following:

Haskell

```
double :: [Double] -> [Double]
double [] = []
double (x:xs) = 2*x : double xs
```

The recursion is done in line 3. $(x:xs)$ is split into its head and its tail separately. What happens is that the head is doubled, and then the function is called on the tail. The head of the first tail is the second element; the head of the second tail is the third element, and so on. Eventually we are left with a blank list, which in this case is the base case/terminating condition. Once the function returns [], all of the doubled elements are added to it through the recursive calls. Calling the function we can see that it works correctly.

Haskell

```
double [1, 2, 3, 4]
Output: [2.0, 4.0, 6.0, 8.0]
```

INSPECTION COPY

COPYRIGHT
PROTECTED



13. Systematic Approach to Problem Solving

This section is all about the skills needed to produce solutions to problems within teams as systems are developed. In addition to being articulate and highly skilled, having strong problem-solving skills are both essential requirements of all modern day computer scientists.

13.1 ASPECTS OF SOFTWARE DEVELOPMENT

Software development is not just the process of coding a piece of software to a specification; it actually encompasses the entire project. This covers *why* you are developing the software, *who* the client is and *what* it needs to do. The most commonly used model of software development is the *waterfall* model where development is divided into specific areas of development, where the result of one area is the input for the succeeding process. Before you can start software development you need a definition of the problem.

- Problem objectives – what the problem is
- Stakeholders – who is affected and has a say in important decisions
- Services and constraints – what the system needs to do and how to do it
- Expectations of ‘quality’ – how the system should function, i.e. ease of use
- Project time frame and budget – how long you have to finish and how much it costs

Software development takes a great deal of time and the most important aspect is communication with the intended users and client. Good communication means that a problem will be clearly defined and a solution to meet those clearly defined objectives is more likely to be accepted.

Prototyping

A prototype is a model of the new system to be developed. This model is the final production of the system is developed. It is used in industry a great deal to help in identifying exactly what the user requires. By using prototypes the user gets a visual representation of what can be made and can make comments before time-consuming mistakes can be made.

There are several methods of prototyping. Two such methods are:

- *Piloting* – using a prototype to test the feasibility of a design proposal
- *Modelling* – built to develop a deeper understanding of the user requirements

The above are *throwaway prototypes*, i.e. once they have achieved their purpose they are discarded. The knowledge found. However, sometimes in development *evolutionary prototypes* are used. A prototype is actually the system under development and is a step closer each time a new feature is added.

In your system development, prototyping methods can be used to an advantage. If you have a new idea, instead of using the whole program and trying to solve a particular problem part you need and developing a separate prototyping system that just does what you need, identifying the exact solution. The advantages of this method are:

1. It allows you to concentrate on just the problem itself, rather than being distracted by other parts of the project/system do not affect the functionality of the system.
2. Other parts of the project/system do not affect the functionality of the system. (For example, a faulty validation routine may stop the code from running. Setting up a separate system with validation routines that just does the new feature to fail.) Setting up a separate system with validation routines that just does the new feature to fail. If it goes wrong it must be the feature you are developing.
3. If your new feature fails catastrophically (for example, an update feature corrupts database data) you do not have to retype/redo any more work, as the system is separate.
4. When complete, introducing into the actual system should be relatively straightforward. It is based on the communication between the system and your new feature. You can ‘copy and paste’ straight into your system and then link within the code.

INSPECTION COPY

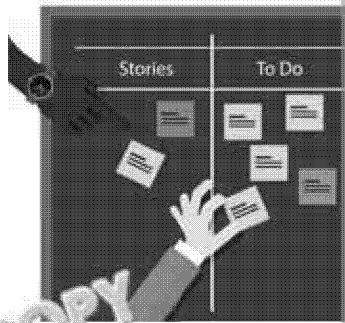
COPYRIGHT
PROTECTED



Agile software development

In larger projects the agile software development management is an approach

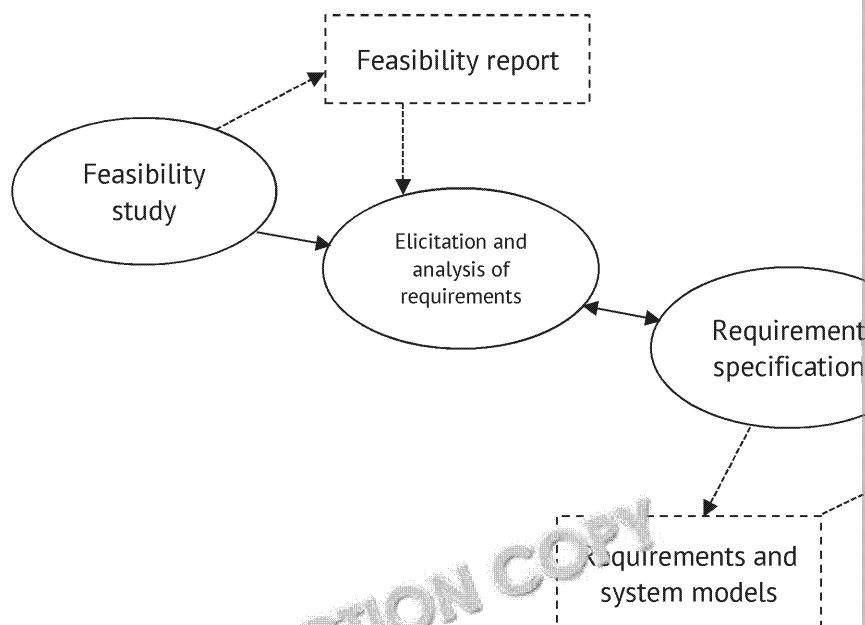
This methodology is to develop the solution through teams. These teams collaborate but focus on particular areas to evolve the project. The solutions produced are used by the client and then actively evolved into a better solution based on feedback. This means that the client sees progress sooner and is able to provide adaptive improvements to allow the project to finally meet its objectives.



ANALYSIS

Analysis of software development is also called *requirement engineering* and is the first stage of the software development process. Before any work can be done, the requirements for the solution must be defined. The system requirements must be defined as well, as these will be what the system must do to function fully and interact with the end user. Two of the most common causes for failure in software development are:

1. Poorly defined requirements – a requirement is poorly defined if it is vague or ambiguous.
2. Poorly managed requirements – if requirements are not managed correctly, the software can keep changing or being added to by the end client.



The *feasibility report* is the examination of whether the system is achievable given the constraints of the requirements.

Requirement elicitation is the method you used to gather the requirements, i.e. interviews with end users, observation of how the current system works and so on.

Requirement analysis is using business scenarios or business models to derive the elicited requirements into their requirement types before the requirements are *specified* to make them precise and detailed.

INSPECTION COPY

COPYRIGHT
PROTECTED



Requirement type	Meaning
Functional	Functional requirements are statements of services that the system must provide and how it should react to certain inputs.
Non-functional	Non-functional requirements are constraints on system performance, usability and efficiency. This also covers security requirements.
Quality	'Quality' can be an intangible property but it covers the system's ability to adhere to standards and the expectations of the system users.
Domain	Domain requirements are the requirements for the system's operating environment, e.g. files being limited to staff with clearance.
Interoperability	These requirements are used for when the new system must be able to use the pre-existing service.
C.R.U.D.	Create, Read, Update and Delete. These requirements are used to specify what data to be handled by the system.

Requirements validation is the final stage and produces the requirements report. Validating requirements is important to make sure that all requirements are realistic. This can be accomplished by running a scenario walkthrough to make sure all requirements are covered, prototyping to find requirements that have been missed during the analysis, and testing a technique that uses a set of conditions to test whether or not a system will meet the requirements.

DESIGN

Using the requirements report from the previous stage you can create what is called a design model. A design model is similar to an architect's blueprint – it is a model of something that you want to build. It allows you to assess the eventual product without building it to test for requirements.

This stage is where the developers look at different aspects of the proposed system. They need to interact for the system to function correctly. Decomposition allows the system to be broken down into *modules* which act as a way of segmenting the full build so that it is more manageable. *Patches*, use modular design to correct mistakes or inefficiencies found in the design. Once the design is complete, database designs are created if the new system needs a database. Finally, pseudocode is created for the developers to produce the working code. This will be used to quickly produce the algorithms that will be robust when the system is built.

The design is validated by checking the proposed design against the requirements of the system, the needs of the user and the 'quality' the client wants. The final design is the input for the next section, *Implementation*.

Design can be developed using the prototyping and agile methods as described in the next section.

**COPYRIGHT
PROTECTED**



IMPLEMENTATION

At this stage, the flesh is put on the bones of the design. This will involve taking the design from the design stage to having it installed for the customer, and the users trained. Steps that need to be possible include:

- Installing software and hardware
- Creating the correct data files
- Properly documenting the system and providing training

This stage can be made easier by the effective use of a *CASE* tool.

CASE (computer-aided software engineering) tools

CASE tools are used to assist in the development of a design into a working

Fourth-generation language This often makes use of concepts such as



This often makes use of concepts such as *code blocks*. The actual machine code they produce may vary, but the nature of such languages is such that the code is tailored to meet the user's requirements.

Interface generators Such as that contained in Visual Basic, which

Such as that contained in Visual Basic, which form's dialogue (menus, textboxes, buttons) write any code.

Code generation facilities To automatically create some source code.

To automatically create some source code (for Applications) systems are written to run Microsoft Word or Excel. Simple operation copying chunks of text can be coded automatically. 'Record', and performs the operation he or system. The operation is translated into code a keyboard shortcut, or assigned to a button code can be created in this way; for example does not allow you to produce iterative or

Data dictionary To record the details of the data in the sys

To record the details of the data in the system in database systems.

Project management tools Software such as *PERT* (Program Evaluation and Review Technique) and *CPM* (Critical Path Method) are used to plan and control the project. These tools help in identifying the critical path and managing resources effectively.

Software such as *PERT* (Program Evaluation with scheduling).

Design tools Such as *desktop publishing (DTP)* packages,

Such as *desktop publishing (DTP)* packages,

Report generator	Used to automatically create documentation
-------------------------	--

Used to automatically create documentation

Implementation in agile development methodology would be through the iterative approach. The most important part (critical path) solution would be deployed before the rest of the solution.

**COPYRIGHT
PROTECTED**



TESTING

Acceptance testing

Acceptance testing is testing to make sure that the system matches the user's requirements, whether the system actually works in practice, and on whether any changes are required. In order for acceptance testing to be useful, the user should be allowed to test the system so that any bugs can be identified, not just given a 'walk-through' that might be prepared. The user must also verify that the interface and 'feel' of the program is right.

Unit testing

Each module in the system is tested to make sure that it functions correctly. (Each individual module or modules) is tested.

Integration testing

The integration test is to make sure that all the units of the system work together. The system must function correctly in themselves, but the whole system must also work. The system must be subjected to:

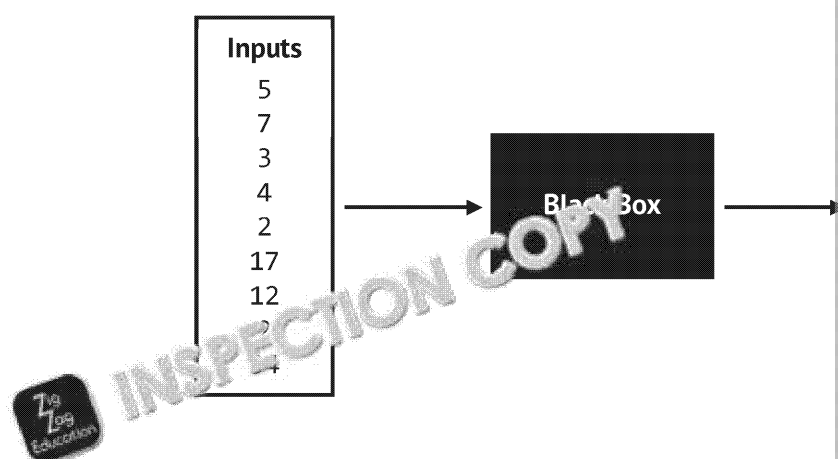
- *Functional tests* to make sure that every aspect of the system functions correctly.
- *Performance tests* to make sure that the system can fulfil its role in a timely manner.
- *Recovery tests* to make sure that the system can recover from various failures.

Black- and white-box testing

These two techniques are aimed at testing the inputs and expected outputs of the system, without knowing the structure of the algorithm. If both these tests are passed, then an assumption is made that all parts of the system are behaving as expected.

In *black-box testing* the program itself is completely ignored, hence the name. The analyst takes a set of inputs with known outputs, put them into the system and compare the actual results with the expected results.

For example, consider the following diagram, where the black box represents a system that is doubling the output:



This system appears to be working correctly so would pass the black-box test.

White-box testing aims to consider all of the possible paths through the algorithm to ensure it is operating as expected. This is done by using the source code being tested to trace the flow of movements through the algorithm. The result of this should be a graph that shows the flow during the design phase. Also, it should be obvious where possible problems are. If there are parts of the graph that you cannot get out of once you have entered and it is impossible to get to, then there is obviously a problem with the code. There are various types of flow diagrams for the white-box testing to avoid having to devote numerous

INSPECTION COPY

COPYRIGHT
PROTECTED



Dry-run testing

A *dry run* is working through a program manually, i.e. on paper. This is used (e.g. it crashes or produces a wrong answer), and so the programmer must work line trying to find out where the problem is. Incidentally, never say in an exam that you go through the code to find problems; always say instead that they will do dry-run testing. This is done using trace tables which are covered in *Section 4.1*.

Choice of test data

A wide range of test data must be entered into the system, in order to test it thoroughly.

- Valid normal data – data of the sort that will be entered into the system (e.g. a valid name)
- Valid boundary data – data that will occur only rarely, but which the system must be able to handle
- Standard incorrect data – this could be data that is only slightly wrong (e.g. a name that is not valid) but is not valid information – to observe the effects on the system of incorrect data
- Standard invalid data – data that should not be entered into the system (e.g. a name that is not valid)
- Extreme data – data that is extreme or unusual

Test plan

Draw up tables containing your test data. There are several areas that you are required to cover in your testing:

- Validation data for all input.
- Data for individual modules (functions and procedures); these may be tested separately from the rest of the system – perhaps at the coding stage.
- Sets of input data – to check each module gives a correct set of output
- Whole-system sets of data (e.g. a day's/week's/month's/year's set of data)

The test plan is usually structured using a table, with the following columns:

- Test number (for later reference)
- Test title (to indicate what is being tested)
- Explanation if necessary of what is being tested
- The test data itself
- The expected result(s) of the test, including where appropriate what the output from the test should be
- The actual result after testing took place – just a yes/no won't do

Justification of test data

You must provide full and detailed justification of why you picked all the test data that you did, and why that is the entire test data you need.

Retesting

Where testing fails (it must), you need to explain what was wrong, what you did to fix it, and when you retested.

If none of your tests failed, there are two possible reasons:

1. There are no errors anywhere in your system (do you really believe that?)
2. You have failed to test your entire system thoroughly enough and have missed the errors.

**COPYRIGHT
PROTECTED**



EVALUATION

Evaluation is carried out for several reasons. It may be carried out by the project team that were drawn up at the analysis stage to check that the program does what was intended by the company purchasing the system before they pay the programmers; agreed criteria. Some organisations will bring in an independent evaluation team to carry out the evaluation.

The evaluation/appraisal may include the following sections:

- Comparison of system against original objectives
- Feedback from actual user
- Improvements needed – how they could be incorporated
- How effective the solution is
- Other possible future developments

Evaluation criteria

The exact criteria by which a project must be assessed will vary from system to system. Generally, the following are the lines for evaluation of a system:

- Whether the system fulfils the user's needs to a satisfactory degree, in terms of function, interface, reliability, efficiency, etc.
- The incidence of system's failure, whether catastrophic or merely irritating
- The amount and cost of maintenance required
- The cost of the system, compared to what was predicted
- The timescale of completion, compared to what was originally projected

It is also important to realise that it may often be necessary to restart or improve the system before you can continue. For example, during the design it may become clear that the system is not understood about the current procedure, so further interviews or questionnaires may be needed. Also, as a result of testing, the whole life cycle may be restarted if it appears that the system is not working. However, provided that the design stage was performed well enough this should not be a problem.

Software maintenance

The evaluation stage may uncover problems in the system, potential for improvement. If the system is not working, it may not mean going back to the problem definition stage and beginning the whole process again. If the changes are needed, rather than a major overhaul, we call this *software maintenance*.

Corrective maintenance

This is where the system has an error that needs to be corrected.

For example, the system may have been tested fully and function correctly by the user. However, the system may fail in a situation that was not expected by the user.

Corrective maintenance is often expensive to both the company and the developer. It is often difficult to trace the cause of the error.

Perfective maintenance

This is where the system functions satisfactorily but improvements to the system are needed.

For example, a screen layout may not present the information in the best way possible. It may be necessary to change a particular area where a change to the system could improve the delay.

INSPECTION COPY

COPYRIGHT
PROTECTED



The above maintenance methods would be performed after a short period of period there would be:

Adaptive maintenance

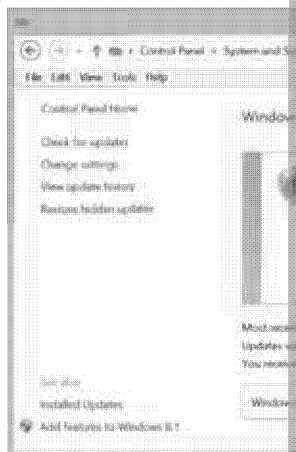
As the organisation expands it may be necessary to alter the system to meet. For example, a single-user system could be adapted to a multi-user system, or a change to the program is necessary (for example, tax changes, or hardware). Sometimes this maintenance may lead to a whole new project and the life cycle.

Modifying systems and patches

The maintenance on systems that have a limited number of users/clients is often done by a team/company.

For systems that are for a larger network it may not be possible for the user to correct all versions directly. Maintenance improvements are often released as programs called patches or service packs. *Patches* are usually for corrective maintenance and *service packs* for perfective/adaptive maintenance.

These programs are run by the user and perform the alterations through a series of instructions (often changing settings, or replacing a program or file).



INSPECTION COPY

COPYRIGHT
PROTECTED



Programming Challenges

Now that you have covered all of the basic theory of programming, you can build on your understanding and build your confidence. Here are six programming challenges that have been devised for you. They start at a basic level and increase in complexity so it's expected that you may find the latter ones particularly challenging. The ★ symbol next to each challenge title indicates the relative difficulty of that challenge.

There are no skills in this section that you have not covered. If you are struggling, refer back to the programming notes in this resource, and break the tasks down into manageable processes. Even if you feel you have a good understanding of programming, it is recommended that you start at the beginning and progress through all tasks.

You should try using as much existing code in all of your solutions where it is possible.

1. TO THE POWER OF...

String manipulation, casting, arithmetic operations

Produce the code that asks the user for two inputs. Without using built-in functions, calculate the power of the first input to the second and return the result to the screen.

2. CONTINUOUS DIVISION

Basic subroutines with parameters, selection, arithmetic operations

Using procedural programming, produce the code that will continuously halve a number until it reaches 1. When the program reaches the number 1 the program should terminate.

Extension:

Modify your code so that it uses a single subroutine that does the division and asks the user for an input and pass the value as a parameter.

3. GUESSING GAME

Random number generation, iteration, subroutines, reading inputs from keyboard

Produce a game where the user has 10 attempts to guess the random number. If the user guesses correctly, the program should say whether the guess was too high or too low.

Extension:

Extend your code so that the user can set the parameters of the game to make it more challenging. The user should be able to set the number of guesses they're allowed and the range of numbers. The program should give feedback to give them a hint as to how close they are to the random number.

INSPECTION COPY

COPYRIGHT
PROTECTED



4. CASE SELECTION

Selection, arithmetic operations, reading inputs from keyboard, subroutines

A hotel needs a new booking system for their rooms. The price of a single night is the culmination of the room type and board basis.

Room type	Board
Single (1): £50	Self-catering: £0
Double (2): £40	Half-board: £10
Family (4): £30	Full board: £20

They also offer a discount for stays over a week. This discount is 25% for every week long. The result should be the total price. The program should show how much has been paid. You should try to break these tasks down into separate subroutines!

5. WRITING TO & READING FROM FILES

Iteration, selection, casting, arithmetic operations, subroutines, arrays, Boolean operations, file input/output

You have been tasked to write three procedures to be used for a leaderboard.

- LoadLeaderBoard() – This *function* should load the values from a file into an array which will then be returned. *You may find it easier to cast the array to a string for ease of use later.*
- PrintLeaderBoard() – This *subroutine* should be passed the leaderboard array and print the values.
- SaveLeaderBoard() – This *subroutine* should be passed the leader board array and save the contents back to the file.

Extension:

Write an additional procedure called CompareScores() which prompts the user to input a score. The program should then iterate through the array from the file to see where the score should appear in the leaderboard. If the score should appear in the leaderboard all other scores after the new score should be input.

6. MAGIC SQUARES

Complex iterations, 2D arrays, Boolean operations, arithmetic operations, date/time

Magic squares are a mathematical phenomenon where all the values in rows and columns add up to a given number. Using a number random generator (between 1 and 10) and a magic square generator where all the numbers sum up to 15. Output the values of the magic square.

For this task you may find it useful to copy each row and column of the array into a string to make it easier.

Note: True magic squares take the diagonals into consideration. For this task you do not need to but you can if you want an extra challenge.

Extension:

Research how to produce a stopwatch timer using the system clock (*hint: System.currentTimeMillis()*). You print the array values the time taken is also printed; you may be surprised at the time taken.

**COPYRIGHT
PROTECTED**



Assembly Programming Challenges

Below is a series of problems to solve. The problems get more and more difficult.

1. *Input two numbers and add them together; output the result.*
2. *Input two numbers and subtract them; output the result.*
3. *Write a program that counts backward from 10.*
4. *Write a counter that counts up to 10.*
5. *Input a number and output its times table up to 10 times.*
6. *Write a counter that counts how many times I enter a non-zero number.*
7. *Write a program which adds up any list of numbers (by looping until I enter a zero).*
8. *Take in two numbers and output them in order (smallest first).*
9. *Take in a number and divide it by 2 (clue: repeated subtraction, counting).*
10. *Input two numbers and find the average (add together and divide by 2).*
11. *Write a program that finds any number entered (a) divided by another number (b) (whole-number part only).*
12. *Write a program that finds the square root of a number entered.*

INSPECTION COPY

COPYRIGHT
PROTECTED



Question Solutions

SECTION 1

Data Types

Question 1

- a) String
- b) Boolean
- c) Float/Double/Real
- d) Integer
- e) Character
- f) Date/Time

Question 2

- a) String – variable must contain multiple characters
- d) Boolean – the account will or will not have an overdraft or not
- b) Float – value must be highly accurate to avoid rounding errors and increase consistency
- e) Date – preferably 'short' date so as not to include time
- c) Integer – only needs to contain a single letter (M or F)
- f) Integer – sort code contains numbers only and is within the size boundary for a signed integer

Programming Concepts

Question 1

- a) Variable Declaration
- b) Variable Declaration
- c) Variable Declaration
- d) Assignment
- e) Assignment
- f) Assignment
- g) Relational
- h) Arithmetic

Question 2

```
If Score > PassBoundary
    OUTPUT PASS
    Select Case gradeCalc
        Score - Pass Boundary >= 30
            Pass = TRUE
            Action ("Grade is A")
        Score - Pass Boundary >= 20
            Pass = TRUE
            Action ("Grade is B")
        Score - Pass Boundary >= 10
            Pass = TRUE
            Action ("Grade is C")
        Score - Pass Boundary >= 0
            Pass = TRUE
            Action ("Grade is D")
    End Select
Else
    Pass = FALSE
End If
End Procedure
```

Pseudocode is structured

- CASE select nested in
- Arithmetic operation for
- must be score minus p
- grade boundaries mus
- correct use of equal to

Set Pass to true and outp

- allow for OUTPUT("Gr

INSPECTION COPY

COPYRIGHT
PROTECTED



Arithmetic Operations

Question 1

- a) $17 \text{ DIV } 8 = 2$
- c) $((16 \text{ DIV } 2) * (6 \text{ MOD } 4)) = 16$
- b) $90 \text{ MOD } 16 = 10$
- d) $26 \text{ MOD } 2 = 0$

Question 2

Var1 \leftarrow Value

Var2 \leftarrow Value

IF Var1 MOD Var2 == 0 Then

 OUTPUT("No remainder!")

Else

 OUTPUT("There is a remainder of " & Var1 MOD Var2)

IF statement used

- only accept the value from user if a

Correct use of MOD

Relative Operations

Question 1

- a) False
- b) True
- c) True

Question 2

Var1 \leftarrow value

Var2 \leftarrow value

Correct use of both greater

Correct outputs for all 3 re

IF Var1 > Var2 Then

 OUTPUT(var1 & "is bigger")

Else If Var1 < Var2 Then

 OUTPUT(Var2 & "is bigger")

Else

 OUTPUT("The numbers are equal")

Boolean Operations

Question 1

- a) False
- b) True
- c) False

String Handling

Question 1

- a) 656.34
- b) 656.34
- c) 11
- d) 8

Question 2

Var1 \leftarrow INPUT

ConvertToInt(Var1)

OUTPUT (Var1 * Var1)

Conversion to integer

Correctly produces s

**COPYRIGHT
PROTECTED**



Random Number Generation

Task

```
isFound ← False
rand ← New Random (1, 10)

WHILE isFound = False
  OUTPUT ("Enter a number between 1 and 10: ")
  userGuess ← INPUT
  IF userGuess == rand Then
    isFound = True
    OUTPUT ("Correct the number was " & userGuess)
  End While
```

Exception Handling

Task

```
Var1 ← Input
Var2 ← Input
Result ← 0

Try
  Var1 ← Convert to Integer
  Var2 ← Convert to Integer
  Result ← Var1 / Var2
  OUTPUT (Result)
Catch Exception
  IF Var2 == 0
    OUTPUT ("You cannot divide by 0")
  END IF
End Try
```

Subroutines

Question 1

Takes the passed value and iterates through the values of 1 to 4 and sums the answer variable minus the number of the loop. This produces a number called a factorial.

Question 2

$3! = 3 * 2 * 1 = 6$

Procedures, Functions and Modules

Question 1

The difference between a function and a subroutine is that a function returns a value to be used to perform arbitrary tasks that do not necessarily require an output.

Question 2

Regardless of the programming language, an error message will be produced either at compile time or runtime which would result in the program failing to compile or crashing.

Question 3

It is considered bad practice because global variables are assigned memory at compile time and the program has closed.

INSPECTION COPY

COPYRIGHT
PROTECTED



Recursive Techniques

Task

OUTPUT MyFib(a)

FUNCTION MyFib (x) # number up to F_x

A ← 0

B ← 1

IF counter >= x then

Return a

ELSE

Temp ← a

B ← temp + a

MyFib(a)

END IF

END FUNC

Q1

- a) The code asks the user for an input and sums all inputs till sum reaches 100.
b) Yes – i*0 is always true but i*0 is always 0 so sum never reaches 100; the escape

Question 2

Returns the address for 'sum' which is continuously stored to stack after each call.

Object-oriented Programming

TASK

Account.Class

CLASS STRUCTURE Account

Private fullName

Private balance

Default constructor

STRUCTURE Account ()

fullName ← "no name"

balance ← 0

END STRUCTURE

partially initialised constructor

STRUCTURE Account (name)

fullName ← Name

balance ← 0

END STRUCTURE

fully initialised constructor

STRUCTURE Account (name, currentBalance)

fullName ← Name

balance ← currentBalance

END STRUCTURE

Class subroutine to an account

PROCEDURE MyDeposit ()

Balance ← balance + 10

OUTPUT deposit

END PROCEDURE

Generator.Class

Declare 3 new c

Account AccountC

Account AccountT

Account AccountT

generate 3 acco

accountOne ← n

accountTwo ← n

accountThree ←

accountThree.My

**COPYRIGHT
PROTECTED**



SECTION 2

Arrays

Question 1

- a) 11
- b) 32

Question 2

- a) [0,3] and [3,0]
- b) [0,1] and [1,0]
- c) By storing [2,3] to a temporary variable and setting [2,3] = [3,3] and [3,3] equal to

Reading and Writing Files

Task

```
Total ← 0
WHILE Input >= -1
    PRINT "Total is: " + total
    PRINT "Enter new integer: "
    inputString ← ReadInput
    total = total + inputString

    fileWriter [fileName]
    writeToFile[inputString]
END WHILE
fileWrite.Close
```

Queues

Question 1

A queue – items are added at one end and removed at the other, meaning the video

Question 2


- a) Head and Tail
 - b) Head == Tail
 - c) Head == Tail + 1 OR Head = 0 Tail = max
 - d) PROCEDURE push(new_item)
 - IF (tail + 1 = head) OR (tail = 10 AND head = 1) THEN
 - PRINT("Queue is full!")
 - ELSE
 - IF tail = 10 THEN
 - tail = 1
 - ELSE
 - tail = tail + 1
 - END IF
 - data(tail) = new_item
 - END IF
- END PROC

INSPECTION COPY

COPYRIGHT
PROTECTED



Question 3



b) Head and tail both equal NULL; both would need to be changed when an item is remove an item, only one needs to be checked for NULL.

c) PROCEDURE push(newItem)

IF head = NULL

head = new element()

tail = head

ELSE

newNode = new element()

head = head.next

END IF

head.next = NULL

head.contents = newItem

END PROCEDURE

Stacks

Question 1

Last In, First Out (LIFO)

Question 2

a) 89, 45, 77, 56

b) FUNCTION add(stack) RETURNS INTEGER

total = 0

WHILE stack is not empty

total = total + stack.pop()

END WHILE

RETURN total

Question 3

a) Billy

b) Check stack is not full, increment top of stack, insert element at top of stack

c) -1 (adding to stack when empty is exactly the same as when not empty)

Graphs

Question 1

a)

Vertex	Connected to
A	C, D, E
B	A, F
C	B, F
D	
E	B
F	B

**COPYRIGHT
PROTECTED**



b)

	A	B	C	D	E	F
A	0	0	1	1	1	0
B	1	0	0	0	0	1
C	0	1	0	0	0	1
D	0	0	0	0	0	0
E	0	1	0	0	0	0
F	0	1	0	0	0	0

c) Adjacency list – it uses less memory resources than using a matrix.

Trees

Question 1

The tree would look like this, with every node being off the left pointer of the one below it.

Question 2
O(n)

Question 3

The simplest way is to add all the parentless items back onto the tree, ignoring their

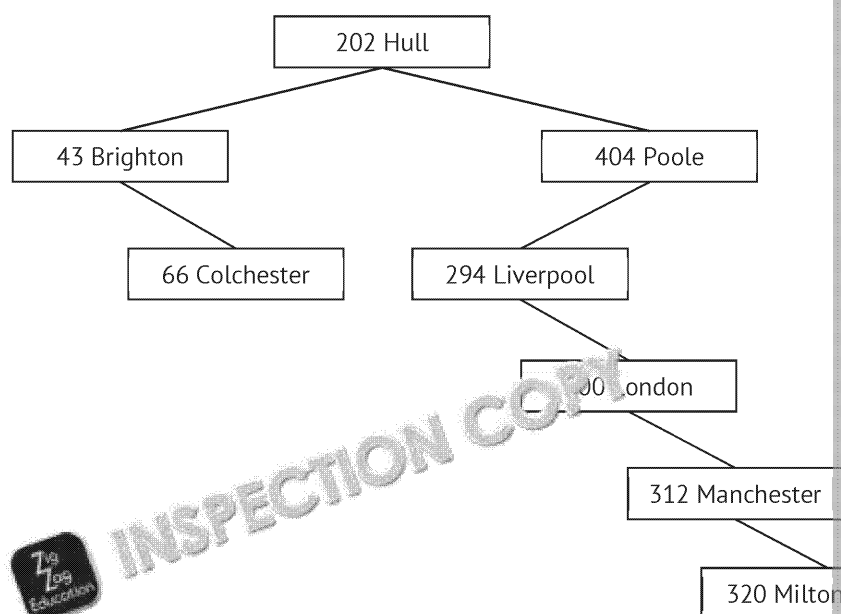
Hash Tables

Question 1

Allows records to be found through an index which can be generated from the index

Question 2

a) + b)



c) Collision

d) A collision resolution strategy – e.g. linked list from each node

**COPYRIGHT
PROTECTED**



Vectors

Question 1

Scaling

Question 2

$$\begin{bmatrix} 3k \\ 0 \end{bmatrix}$$

Question 3

$$A = \begin{bmatrix} 3 \\ 1 \end{bmatrix} \quad C = \begin{bmatrix} 4 \\ 3 \end{bmatrix} \quad \text{Therefore, } A + C = \begin{bmatrix} 3 \\ 1 \end{bmatrix} + \begin{bmatrix} 4 \\ 3 \end{bmatrix} = \begin{bmatrix} 7 \\ 4 \end{bmatrix}$$

SECTION 3

Graph Traversal

Q1

- a) 1, 3, 5, 2, 4, 7, 6
- b) 1, 3, 5, 2, 4, 6, 7

Tree Traversal

Question 1

- a) 0, -1, -2, -5, -7, -6, -4, -3, 1, 2, 5, 3, 4, 6, 7
- b) -7, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 7
- c) -6, -7, -3, -4, -5, -2, -1, 0, 4, 3, 7, 6, 5, 2, 1

Reverse Polish Notation

Question 1

RPN removes the need for brackets as computers do not understand how to use them

Question 2

- a) 56 / 5 +
- b) 77 * 625++
- c) 457/06-+

Question 3

- a) a*b
- b) g + h - b
- c) b/m + g/h

Searching Algorithms

Q1

- a) A binary search cannot be performed because the data items are not sorted into a

- b)

1	2	3	4	6	8	9
---	---	---	---	---	---	---

It would take two passes to find the number 2.

**COPYRIGHT
PROTECTED**



Question 2

Index	1	2	3	4	5	6	7	8	9	10
Data	A	C	E	J	L	O	Q	R	Y	Z
	L				M				R	

Index	1	2	3	4	5	6	7	8	9	10
Data	A	C	E	J	L	O	Q	R	Y	Z
						L		M		R

R is found in index 8.

Shortest Path Algorithm

Question 1

Shortest path (least cost):

$$\vec{ac} = 1$$

$$\vec{cb} = 2$$

$$\vec{be} = 3$$

$$\vec{ed} = 1$$

$$\vec{df} = 3$$

$$\vec{fg} = 4$$

$$\vec{gh} = 3$$

$$\vec{ah} = 17$$

a, c, b, e, d, f, g, h

SECTION 4

Finite-state Machines

Question 1

- Input alphabet = {open door, close door, timeout}
- Output alphabet = {light on, light off, alarm on, alarm off}
-

Current State	Input	Next State	Output
S1	open door	S2	light on
S2	close door	S1	light off
S2	timeout	S3	alarm on
S3	close door	S1	alarm off, light off

Question 2

Current State	Input	Next State	Output
S1	sensor triggered	S2	open doors, reset timer
S2	sensor triggered	S1	reset timer
S2	timeout	S3	close doors

COPYRIGHT
PROTECTED



Regular Languages and Expression Notation

Question 1

Intersection can be applied to check what numbers appear in both and they are: 6 and

Question 2



2, 3, 5, 10

Question 3

- a) $^*[a-z]$
- b) False – all letters of the English language appear in the sentence

Context-free Languages

Question 1

- a) False
- b) False
- c) 
- d) 

Classification of Algorithms

Question 1

$(n+1) + n + (n+1) + n = 4n + 2$

Question 2

- a) Polynomial
- b) Exponential
- c) Linear

Question 3

Its order of complexity is calculated from the worst-case scenario run-time.

Question 4


An algorithm with a worse space complexity manages system resources poorly and w

A Model of Computation

Question 1

- a) Turing machines are still used to this day because they allow the study of what is modern computers.
- b) Modern computers operate on a similar principle: programs are stored in the same a universal Turing machine.
- c) No – it also has a 'state of mind' which defines what it should do when given a p

Question 2

- a) ight
- iii) Yes
- b) The Turing machine enters into an infinite loop.
- c) It will add two numbers together.

**COPYRIGHT
PROTECTED**



SECTION 5

Number Systems

Question 1

Natural numbers are all integers without negatives, meaning they can all be used in counting numbers; whereas real numbers encompass all the types of sets including a large amount of memory due to their value of the accuracy.

Number Bases

Question 1

- a) $26_{10} \rightarrow 00011000_2$
- b) $100_{10} \rightarrow 01100100_2 \rightarrow 64_{16}$
- c) $7C_{16} \rightarrow 1111010_2 \rightarrow 122_{10}$
- d) $01001001_2 \rightarrow 73_{10}$
- e) $18_8 \rightarrow 1001100_2 \rightarrow BC_{16}$
- f) $201_{10} \rightarrow 11010011_2 \rightarrow 201_{10}$

Question 2

The highest value that can be stored in a single byte is 256 (2^8). In order to store anything larger than 256, you need more than one byte.

Units of Information

Question 1

512 bits

Question 2

$$2^{32} = 4,294,967,296$$

(four billion, two hundred and ninety-four thousand, nine hundred and sixty-seven thousand, two hundred and ninety-six)

Question 3

$$2^{40} = 1.0995116 \times 10^{12} \text{ or } 1,099,511,600,000 \text{ (one trillion, ninety-nine billion, five hundred and ninety-five million, one hundred and thirty-six thousand, nine hundred and fifty-two)}$$

Binary Number Systems

Question 1

- a) 00111100_2
- b) 00111000_2
- c) 10001001_2
- d) 01011010_2

Question 2

- a) 00100011_2
- b) 00001100_2

Question 3

No. of bits in a byte is 8. The value is -128 (10000000_2).

Information Coding Systems

Question 1

1

Question 2

6

Question 3

000 111 000 111 000 000 000 111

INSPECTION COPY

COPYRIGHT
PROTECTED



Bitmapped Graphics

Question 1

$$7 * 7 = 49 \text{ inches}^2$$

$$80^2 = 6400 \text{ pixels per inch}$$

$$6400 * 49 = 313,600 \text{ pixels in total}$$

4 bits required to store the 8 colours

$$\text{File size} = 313,600 * 4 = \left(\frac{1,254,400}{8000}\right) = 156 \text{ kilobytes}$$

Representing Sound

Question 1

$$(8000 * 16) * 30 = \left(\frac{3,840,000}{8000}\right) = 480 \text{ kilobytes}$$

Question 2

$$\text{Sample rate} = \left(\frac{\text{file size / bits}}{\text{length}}\right) = \left(\frac{\frac{100,000}{10}}{10}\right) = 1,000 \text{ Hz}$$

Question 3

Because no matter how many bits you have to represent the value of the pulse peaks be lost which leads to the staircase effect.

Data Compression and Encryption

Question 1

Lossy compression identifies seemingly redundant data and removes it from the file.

Question 2

'Computer Science'

Question 3

It needs to be at least the same length because otherwise the modulo division no longer works. If the key is longer it just means that part of the key is not used.

Question 4

Yes – the cipher effectively works as a one-time key. As the keys are never reused, it is secure. If the key was shared it would always be impervious to all attempts to break the key values.

SECTION 6

Hardware and software

Question 1

- a) Utility software
- b) Operating system
- c) Translation program
- d) Linker program

Question 2

Abstraction is created by the operating acting as a virtual machine where the complexity is hidden from them.

Question 3

The main drawback is that the code is not checked while it is being programmed. It is only when the program is run that errors are found.

**COPYRIGHT
PROTECTED**



Classification of Programming Languages

Question 1

- a) The code reads the passed value and compares it to 0 – if the value is not 0 then the value of sign is then returned.
If the value is less than 0, then it is set to -1; if it is greater than 0 then the value of sign is then returned.
- b) The output would be 1.

Question 2

Assembly language

Uses mnemonics
Some abstraction
Uses assembler
Easier to program in
Runs more slowly due to translation

Machine language

Comprised entirely of 1s and 0s or hexade
No abstraction
Doesn't need translating
Harder to program in
Runs faster

Question 3

Yes – although all compilers are different, the high-level languages are more portable as the translation is automated by the computer. As long as the compiler can understand the language, the program will compile on any computer.

Logic Gates

Question 1

- a) True b) False c) False


Question 2

Input A	Input B	Input C	Input D	Output
0	0	0	1	0
0	0	1	0	0
0	1	0	0	1
1	0	0	0	1
1	1	1	1	0
1	1	1	0	1
1	1	0	1	1
1	0	1	1	0
0	1	1	1	0
1	1	0	0	1
1	0	0	1	1
0	0	1	1	0
0	1	1	0	1
0	1	0	1	1
1	0	1	0	1
0	0	0	0	0

1 = true, 0 = false



Q



COPYRIGHT
PROTECTED



Boolean Algebra

Question 1

- a) $A + (A \cdot B) = A$
- b) $A \cdot B + A \cdot \bar{B} = A$
- c) $(A + A \cdot B) + (A + A \cdot B) = A$
- d) $\overline{A + B} = \overline{A \cdot B}$
- e) $\overline{(A + B)} + B = B$

SECTION 7

Internal hardware components of a computer

Question 1

The address bus carries the value of the memory address being indexed. If a valid memory address is received, then data is permitted to be conveyed via the data bus. The flow of data is controlled by the control bus.

Question 2

RAM stands for Random Access Memory. It is used by the system as a storage medium for current tasks, used by the operating system and used to store values during run-time. After current is lost from the board all data that was stored within the memory module is lost.

Question 3

- a) Motherboard
- b) Cache memory
- c) CPU

Structure and the Role of the Processor

Question 1

Arithmetic logic unit (ALU)

Question 2

The contents of the program counter are transferred to the memory address buffer, which then transfers the information into the memory buffer register (or memory data register). The information is then transferred to the ALU. This is called the fetch part or phase of the cycle.

The control unit then splits the information into two parts: the operation code and the operand. The operation code tells the computer which instruction to perform. The control unit then switches the processor to perform the action. This is called the decode phase.

The processor then transfers the data (operand) to the appropriate part of the processor. The operand would be transferred to the MAR and used appropriately, otherwise it would be discarded. The final phase is called the execute phase.

Question 3

- a) The computer is constantly running background tasks/services that ensure the computer is up to date. It connects to time servers to update the time allowing you to connect to the internet. It is constantly being exchanged.
- b) Yes – an interrupt is sent to the CPU; the need to load an application is of higher priority than the current task undergone while the computer is idle. Once the interrupt is serviced and the application is loaded, the computer can then continue with other background tasks.

Question 4

- a) False
- b) True
- c) False
- d) True

**COPYRIGHT
PROTECTED**



SECTION 8

Computing Issues

Question 1

A company must consider that constant monitoring of its employees may be considered intrusive to the employees. The employees themselves may feel that this monitoring is intrusive to their privacy. They may also feel resentful towards this practice. Another aspect the company must consider is that if the call is monitored it is possible that employees may give a poorer service to customers. If the quality of the call is monitored it is possible that employees may give a poorer service to customers.

Question 2

Pirated software is stealing software, and just like any other product or service the market is reduced. People that have created the software. Financially, the piracy industry removes a large part of the market. Consequently the prices may have to be increased to cover their costs.

Equally, the quality of the software is reduced as pirates will often alter software to be able to use it without the need for a license. These altered or 'cracked' programs often contain mistakes and are not as robust.

Question 3

Often government agencies will use data to track and monitor individuals in society. This is done for national security. However with the use of drones over highly populated areas, the government is monitoring individual rather than suspects. The 'big brother' effect is that the government would be monitoring everyone. Many would consider this an invasion of privacy and be concerned about how the data is used.

Question 4

People are spending more time stationary and performing repetitive actions.

RSI – Repetative strain injury occurs with people that spend a lot of time typing on a computer. It is also common in people that performed repetitive tasks in factories.

Eye strain is becoming more of a problem with computers which previously was not common. People are in environments in which they would not focus on a fixed point for many hours (computer work).

Back ache has issues in common with industrial accidents in which people were required to stand for long periods (accountants/book keeping). Mobile devices such as laptops make it difficult to be seen.

Stress has increased with the use of computers as information is required quicker and more often. In industries 50 years ago. A computer failure in modern age can be life threatening and information often suffer from stress.

SECTION 9

Communication

Question 1

Serial communication can operate either in synchronous mode (with a mutual clock) or asynchronous mode.

Question 2

- a) False
- b) False
- c) True

Question 3

Control bus

INSPECTION COPY

COPYRIGHT
PROTECTED



Networking

Question 1

No performance degradation

Secure communication – communication is direct therefore no eavesdropping is present

Scalable – easy to add new nodes

Question 2

Wireless network adapter

Question 3

Method

Application

Hiding the network SSID

Keeps network hidden and secure to all but those who know the SSID

MAC address white listing

Even if someone does manage to gain knowledge of your network, white listing prevents all but those on the list from connecting

Wi-Fi Protected Access

Prevents people connecting to the network if they don't have the correct packets being sent that bluff the passphrase to gain access

The Internet

Question 1

The payload is the volume of data that is being transmitted – the useful part of the packet

Question 2

Packet switching is the digital equivalent of a switchboard that allows packets to be sent. The route the packets take is completely random and independent of any other packets being sent.

Question 3

A firewall that operates at the packet level and studies each packet as it arrives and passes it to the router; each packet must pass the packet policy to gain access to the network.

Question 4

Computer viruses tend to copy themselves, infect the core system files and hide themselves from any given moment. Trojans are not self-replicating and tend to be hidden among innocent files. They are active when the file is executed. They are used to gain access into a system and take control.

Transmission Control Protocol / Internet Protocol

Question 1

- An IP address stands for Internet Protocol address and acts as the address of the computer. It is used to transfer data to and from the computer.
- Also known as routable and non-routable, these are the different variations of IP addresses. A public IP address will have a public IP address that is routable – that is, data can be sent to and from it. A private IP address will have a private IP address which is non-routable that the router uses to connect to the internet.
- A subnet mask is used to split an IP address into the host identifier and the network identifier.

Question 2

The Media Access Control (MAC) address is used to identify a device and to communicate with each device.

Question 3

As it uses 128 bits the IP address pool becomes much larger; it could remove the need for NAT. Many addresses available it would be near impossible to use them all in the foreseeable future.

Question 4

NAT is the 'workaround' used by IPv4-enabled routers to increase the number of available public addresses. NAT allows the router to apply a private network address to each device and then forward traffic to that device only.

**COPYRIGHT
PROTECTED**



SECTION 10

Conceptual Data Models and Entity Relationship Modelling

Question 1

The models are used as a way of gaining information about how data flows in a system.

Question 2

e.g. Salary(SalaryCode, SalaryValue)

Relational Databases

Question 1

This is where the data is stored in different tables which are linked together using primary and foreign keys.

Question 2

- a) EmployeeID. A key is an item of data which is a unique identifier for that particular entity. In this case, the EmployeeID.
- b) SalaryCode. A foreign key is a primary key in another table/entity within the database. It links the salary entity to the employee entity. In this case the employee will be linked to the salary entity.

Database Design and Normalisation

Question 1

- a) Separate any attributes from keys formed in the previous step that are only dependent on the primary key.
- b) Separate any attributes that are dependent on other non-key attributes; foreign keys. If a part of a key can be derived from other attributes, then it is a composite key for redundant parts. If a part of a key can be derived from other attributes, then it is a composite key for redundant parts.

Question 2

Normalisation increases speeds of data retrieval by structuring data in an easy-to-understand way because of this structure as there is less likely to be anomalous data in the fields.

SECTION 11

Big Data

Question 1

Big data is a term given to any data source that is difficult to process due to its lack of structure, volume, variety or velocity.

Question 2

By allowing programmers to produce code more efficiently.

- Immutable data structures
- Stateless functions
- High-order functions

INSPECTION COPY

COPYRIGHT
PROTECTED



PROGRAMMING CHALLENGES

1. To the Power Of...

```
SUBROUTINE RunProgram ()
    Output "enter the first value: "
    baseVal <-- UserInput
    Output "enter the second value: "
    powerVal <-- UserInput
    Output baseVal^powerVal = ToThePower(baseVal, powerVal)
END SUBROUTINE

FUNCTION ToThePower (Base, Power)
    temp <-- 0
    For i <-- 1 to power
        temp + = Base
    For
    n temp
END FUNCTION
```

2. Continuous Division

```
PROCEDURE RunProgram ( )
    OUTPUT "Enter an integer number: "
    userInput ← input # Converted to integer
    StartDivision (userInput)
END PROCEDURE

PROCEDURE StartDivision ( n )
    WHILE n > 1
        n ← n / 2
        OUTPUT n
    End While
END PROCEDURE
```

3. Guessing Game

```
PROCEDURE Initialise ( )
    OUTPUT "Enter the lowest value: "
    lowest ← Input
    OUTPUT "Enter the highest value: "
    highest ← Input
    secretNumber ← CreateRandom(lowest, highest)
    OUTPUT "How many guesses would you like to have?"
    guesses ← Input
    PlayGame(secretNumber, guesses)
END PROCEDURE

FUNCTION CreateRandom (minimum, maximum)
    tempRandom ← 0
    tempRandom ← NewRandom(minimum, maximum)
    return tempRandom
END FUNCTION
```

INSPECTION COPY

COPYRIGHT
PROTECTED



```

PROCEDURE PlayGame (secret, maxGuesses)
    isFound ← False
    totalGuesses ← 0

    While isFound ← False AND totalGuesses < maxGuesses
        OUTPUT "Enter your guess: "
        playerGuess ← input
        totalGuess++

        If playerGuess ← secret
            OUTPUT "Congratulations! You've won in " & totalGuesses
            OUTPUT "Would you like to play again?"

            If replay = y Then
                isFound ← false
                If replay = n Then
                    isFound ← true
                    break
            End If
        Else if playerGuess - secret < 5 AND playerGuess - secret > -5 Then
            OUTPUT "So close!"
        Else if playerGuess - secret < 10 AND playerGuess - secret > 10 Then
            OUTPUT "Quite close!"
        Else
            OUTPUT "Not even close"
        End If
    End While
END PROCEDURE

```

4. Case Selection

```

PROCEDURE RunProgram ( )
    newRoom ← y
    While newRoom ← y
        NewBooking( )
        OUTPUT "Would you like to calculate a new room?"
        newRoom ← Input
    End While
END PROCEDURE

PROCEDURE NewBooking( )
    input ← 0
    OUTPUT "1 for single, 2 for twin, 3 for double, 4 for family"
    Total ← room(input)
    OUTPUT "1 for self-catered, 2 for half-board, 3 for full board"
    boardTotal ← board(input)
    OUTPUT "Total: " & totalCost(roomTotal, boardTotal)
END PROCEDURE

```

...

**COPYRIGHT
PROTECTED**



...

PROCEDURE Room (input)

tempTotal \leftarrow 0

switch (input)

case 1:

tempTotal \leftarrow 50

break

case 2:

tempTotal \leftarrow 75

break

case 3:

tempTotal \leftarrow 90

break

return tempTotal

END PROCEDURE

PROCEDURE Board (input)

tempTotal \leftarrow 0

switch (input)

case 1:

tempTotal \leftarrow 0

break

case 2:

tempTotal \leftarrow 5

break

case 3:

tempTotal \leftarrow 10

break

return tempTotal

END PROCEDURE

FUNCTION TotalCost (Room, Board)

OUTPUT "How long would you like to stay for?"

stayLength \leftarrow input

roomCost \leftarrow Room + Board

tempTotal \leftarrow 0

If stayLength > 7 Then

tempTotal \leftarrow (roomCost * 7) - ((stayLength - 7) * (roomCost * 0.25))

Else

tempTotal \leftarrow roomTotal * stayLength

End if

Return tempTotal

END FUNCTION

5. Writing and Reading from Files

PROCEDURE RunProgram ()

newArray [] \leftarrow LoadLeaderBoard ()

PrintArray(newArray)

OUTPUT "Enter new score: "

newInput \leftarrow Input

compareScore(newArray, newInput)

END PROCEDURE

...

**COPYRIGHT
PROTECTED**




```

...
PROCEDURE PrintArray (arrayName[] )
    OUTPUT "Current leaderboard is: "
    For i ← 0 to array.Length
        OUTPUT (array[i] + " ")
    End For
END PROCEDURE

FUNCTION LoadLeaderBoard ( )
    leaderboardArray[] ← leaderboardArray [5]
    lineFromFile ← Null
    new StreamReader ← reader
    fileName ← #newFileName
    reader ← new streamreader (fileName)
    counter ← 0
    while (!reader.EndOfStream)
        lineFromFile ← reader.ReadLine()
        leaderboardArray[counter] ← lineFromFile
        counter++
    End While
    Close Reader
    Return leaderboardArray
END PROCEDURE

PROCEDURE SaveLeaderBoard (leaderboard[])
    StreamWriter ← fileWriter
    fileName ← #newFileName
    FileWriter ← new FileWriter(fileName)
    for i ← 0 to leaderboard.Length
        inputFromArray ← leaderboard[i]
        FileWriter.WriteLine(inputFromArray)
    End For
    OUTPUT "New entry added. Leaderboard saved."
    FileWriter.Close()
END PROCEDURE

PROCEDURE CompareScore (newArray[], newScore)
    temp1, temp2 ← 0
    If newScore < newArray[4]
        OUTPUT "Not on leaderboard"
    Else
        For i ← 0 to newArray.Length
            If newScore > newArray[i]
                temp1 ← newArray[i]
                newArray[i] ← newScore
                For j ← i + 1 to newArray.Length
                    temp2 ← newArray[j]
                    newArray[j] ← temp1
                    temp1 ← temp2
                End For
            End If
        End For
        PrintArray(newArray)
        SaveLeaderBoard(newArray)
        Break
    End For
End If
END PROCEDURE

```

**COPYRIGHT
PROTECTED**



6. Magic Squares

```
SUBROUTINE RunProgram ()
    matrixIsMagix <-- false
    counter <-- 1
    while matrixIsMagic == false
        theMatrix[3,3]
        theMatrix <-- PopulateArray(theMatrix)
        rowsAreMagic <-- PerformRowCheck(theMatrix)
        colsAreMagic <-- PerformColCheck(theMatrix)
        OUPUT "Squares produced: " + Counter
        counter++
        if rowsAreMagic && colsAreMagix == true
            output "magic square for " + Counter
            PrintArray(theMatrix)
            break
        Else
            While
    END SUBROUTINE
```

```
FUNCTION PopulateArray (tempArray)
    rand <-- New Random
    for i <-- 1 to 3
        for j <-- 1 to 3
            tempArray[i,j] <-- rand.Next(1 to 10)
        End For
    End For
    return tempArray
END FUNCTION
```

```
SUBROUTINE PrintArray (tempArray)
    for i <-- 1 to 3
        for j <-- 1 to 3
            output tempArray[i,j]
        End For
        Output " " # Print to new line
    End For
END SUBROUTINE
```

```
FUNCTION PerformRowCheck (tempArray)
    isMagic <-- false
    row[3]
    for i <-- 0 to 3
        for j <-- 0 to 3
            row[i] <-- tempArray[i,j]
            tempTotal <-- tempTotal + row[j]
        End For
        If tempTotal != 15
            break
        End if
    End For
    End For
    return isMagic
END FUNCTION
...
```

INSPECTION COPY

**COPYRIGHT
PROTECTED**



...

```
FUNCTION PerformColCheck (tempArray)
  isMagic <-- false
  col[3]
  for i <-- 0 to 3
    for j <-- 0 to 3
      col[j] <-- tempArray[j,i]
      tempTotal <-- tempTotal + col[j]

      If tempTotal != 15
        break
      End if
    End For
  End For
  return isMagic
END FUNCTION
```



ASSEMBLY PROGRAMMING CHALLENGES

Solutions below are for Little Man Computer (LMC).

1. Input two numbers and add them together; output the result.

```
INP
STA A
INP
ADD A
OUT
HLT
A DAT
```

2. Input two numbers and subtract them; output the result.

```
INP
STA A
INP
SUB A
OUT
HLT
A DAT
```

3. Write a program that counts backwards from 10.

```
START OUT
SUB ONE
BRZ END
BRA START
END HLT
COUNT DAT 10
ONE DAT 1
```



INSPECTION COPY


**COPYRIGHT
PROTECTED**



4. Write a counter that counts up to 10.

```
START LDA COUNT
      ADD ONE
      STA COUNT
      OUT
      SUB TEN
      BRZ END
BRA START
END   HLT
COUNT DAT
ONE DAT 1
TEN DAT 10
```


5. Input a number and output its times table up to 10 times.



```
INP A
LDA TOTAL
ADD A
STA TOTAL
OUT
LDA COUNT
ADD ONE
STA COUNT
SUB TEN
BRZ END
BRA START
END   HLT
A DAT
TOTAL DAT
COUNT DAT
ONE DAT 1
TEN DAT 10
```

There are different methods to do this; however, above is probably the easiest to remember count at set-up is 0 not 1.

6. Write a counter that counts how many times I enter a non-zero number.



```
START INP
      BRZ END
      LDA COUNT
      ADD ONE
      STA COUNT
      INP A
      LDA COUNT
      OUT
      HLT
COUNT DAT
ONE DAT 1
```

This introduces the idea that we can branch out at any point and then load the i

**COPYRIGHT
PROTECTED**



7. Write a program which adds up any list of numbers (by looping until I enter numbers I entered).

```

START INP
      BRZ END
      ADD TOTAL
      STA TOTAL
      LDA COUNT
      ADD ONE
      STA COUNT
      BRA START
END   LDA TOTAL
      OUT
      LDA COUNT
      OUT
      HLT
TOTAL DAT 0
COUNT DAT 1

```

8. Take in two numbers and output them in order (smallest first).

```

INP
STA A
INP
STA B
SUB A
BRP RESULT
LDA B
STA TEMP
LDA A
STA B
LDA TEMP
STA A
RESULT   LDA A
OUT
LDA B
OUT
HLT
A DAT
B DAT
TEMP DAT

```

This introduces the idea of using a temp location to allow the values to swap over

**COPYRIGHT
PROTECTED**



9. Take in a number and divide it by 2 (clue: repeated subtraction, counting

```
INP
START SUB TWO
    BRP GOES
    BRA END
GOES STA A
    LDA COUNT
    ADD ONE
    STA COUNT
    LDA A
BRA START
END LDA COUNT
OUT
HLT
A DAT
TWO DAT 2
COUNT DAT 1
ONE DAT 1
```

This gives the idea that you jump out to run a routine then jump back in, i.e. I test the count and then jump back in.

10. Input two numbers and find the average (add together and divide by 2).

```
INP
STA A
INP
ADD A
START SUB TWO
    BRP GOES
    BRA END
GOES STA A
    LDA COUNT
    ADD ONE
    STA COUNT
    LDA A
BRA START
END LDA COUNT
OUT
HLT
A DAT
TWO DAT 2
COUNT DAT 1
ONE DAT 1
```

Notice simple bit at the top; just take in a number, take in another and add

INSPECTION COPY

COPYRIGHT
PROTECTED



11. Write a program that finds any number entered (a) divided by another number

```

    INP
    STA A
    INP
    STA B
    LDA A
    START SUB B
        BRP GOES
        BRA END
    GOES STA A
        LDA COUNT
        ADD ONE
        STA COUNT
        LDA A
        BRA START
    END LDA COUNT
    A DAT
    B DAT 4
    COUNT DAT
    ONE DAT 1

```

12. Write a program that finds the square root of a number entered.

```

    INP
    STA NUMBER
    LOOP SUB MINUS
    STA NUMBER
    LDA MINUS
    ADD INCREASE
    STA MINUS
    LDA COUNT
    ADD INC
    STA COUNT
    LDA NUMBER
    BRZ END
    BRP LOOP
    BRA UNABLE
    END LDA COUNT
    OUT COUNT
    HLT
    UNABLE LDA MINUS
    OUT MINUS
    HLT
    NUMBER DAT
    MINUS DAT 1
    INCREASE DAT 2
    COUNT DAT
    INC DAT 1
    FAIL DAT 0

```

INSPECTION COPY

**COPYRIGHT
PROTECTED**

