

FAULT MANAGEMENT SYSTEM
FOR A TELECOMMUNICATIONS NETWORK

This invention relates to a fault management system for managing faults in
5 the terminating circuits of a telecommunications network and also to a method of
operating such a fault management system.

A conventional public telecommunications network comprises a relatively
small number of interconnected main switches and a much larger number of local
switches, each of which is connected to one or two main switches. The local
10 switches are connected to the terminating circuits of the network and the far ends
of these circuits are connected to terminal equipment such as telephone
instruments provided for users of the network. The network formed from the main
switches and local switches is known as the core network while a network formed
from the terminating circuits is known variously as an access network or a local
15 loop. In this specification, it will be referred to as an access network. Some
terminating circuits are connected to a remote concentrator, which may or may not
have switching capability. The remote concentrator is then connected to a local
switch. In this specification, the term "local switch" is to be interpreted to cover
both local switches and remote concentrators.

20 In a conventional access network, each terminating circuit is formed from
a pair of copper wires. Typically, each pair of copper wires pass through a series
of nodes between the local switch and terminal equipment. Examples of such
nodes are primary cross-connect points, secondary cross-connect points,
distribution points and joints.

25 Recently, optical fibres have been used to carry terminating circuits in
access networks. In a modern access network, both pairs of copper wires and
optical fibres are used to carry the terminating circuits. Where a terminating circuit
is carried by an optical fibre, the circuit will typically pass through several nodes
between the local switch and the terminal equipment. At each node, the incoming
30 fibre from the local switch is split into a group of outgoing fibres which branch out
in various directions. Where a terminating circuit is carried by an optical fibre
from the local switch, the last part of the circuit may be carried by a pair of copper
wires.

Unfortunately, terminating circuits are prone to faults. In a conventional access network, examples of such faults are disconnection, a short circuit between the two wires of a pair of wires, and a short circuit between one of the wires and earth. The causes of the faults include ingress of water into a node and
5 also physical damage to a node.

Local switches are provided with circuit testing apparatus which may be used to test its terminating circuits. When a customer reports a fault on a terminating circuit, the circuit may then be tested to identify the fault condition. In order to repair the fault, it is present practice for an engineer to make a guess at
10 the location of the fault and open the node at that location. If the engineer's guess is wrong, he will have to open one or more further nodes before finding and repairing the fault. Sometimes an engineer will have to open each fault in turn between the local switch and the terminal equipment until he finds a node where the fault is present. Typically, in a conventional access network, an engineer will
15 open between 2.5 and 3 nodes on average before locating a fault.

The present practice for locating faults thus suffers from two problems. Firstly, it is time consuming for an engineer to have to open several nodes before locating and repairing a fault. Secondly, as the nodes are of a delicate construction, each time an engineer opens a node he will damage the node with
20 the result that there is another fault on a terminating circuit.

According to one aspect of this invention there is provided a fault management system for a telecommunications network including a local switch and a set of terminating circuits extending between said local switch and terminal equipment provided for users of the network, each of said terminating circuits
25 passing through a series of nodes between said local switch and its respective terminal equipment, said fault management system comprising:

circuit testing apparatus located at said local switch and arranged to perform circuit tests on said terminating circuits;

a store containing data relating to said terminating circuits and said nodes;

30 means for instructing said circuit testing apparatus to perform a set of tests on one of said terminating circuits;

This Sheet Has Been Replaced



THIS SHEET HAS BEEN
REPLACED

means for checking the results of a set of tests performed by said circuit testing apparatus for the presence of a suspected fault, said checking means being arranged to produce a fault report when a fault is suspected;

means for identifying the nodes of a terminating circuit on which a
5 suspected fault is present;

means for evaluating a score for each node on a terminating circuit line on which a fault is suspected which represents the likelihood of the suspected fault being present at the node, said evaluating means using the fault report relating to the suspected fault and data contained in said data store in evaluating the score;
10 and

means for ranking nodes of a terminating circuit on which a fault is suspected in accordance with their scores found by said evaluating means;

whereby, in use, after testing a circuit on which a fault is suspected, a list is produced of the nodes on said circuit in which the nodes are ranked in
15 accordance with the likelihood of the fault being present at each node.

With this invention, as a result of ranking the nodes on a terminating circuit on which a fault is suspected in accordance with the likelihood of the fault being present at each node, an engineer can be directed to the node at which it is most likely that the fault is present. Although with this invention the engineer is
20 not directed to the correct node on every occasion, the present invention does reduce the number of nodes which an engineer has to open before locating faults. On average, when using the invention in a conventional access network, it is estimated that an engineer will need to open less than 1.5 nodes before finding the node at which a fault is present.

25 According to a second aspect of this invention, there is provided a fault management system for a telecommunications network, said telecommunications network including a local switch and a set of terminating circuits extending between said local switch and terminal equipment provided for users of the network, each of said terminating circuits passing through a series of nodes
30 between said local switch and its respective terminal equipment, said fault management system comprising:

circuit testing apparatus located at said local switch and arranged to perform circuit tests on said terminating circuits;

a computer system for operating said circuit testing apparatus;

said computer system including a store containing data relating to said terminating circuits and said nodes;

said computer system being controlled by at least one program so as to
5 perform the following operations:

instruct said circuit testing apparatus to perform a set of tests on one of
said terminating circuits;

check the results of a set of tests performed by said circuit testing
apparatus for the presence of a suspected fault and produce a fault report when a
10 suspected fault is present;

identify the nodes of a terminating circuit on which a suspected fault is
present;

evaluate a score for each node on a terminating circuit on which a fault is
suspected which represents the likelihood of the suspected fault being present at
15 the node by using the fault report relating to the suspected fault and data
contained in said store; and

rank nodes of a terminating circuit on which a fault is suspected in
accordance with their scores.

According to a third aspect of this invention, there is provided a method of
20 operating a fault management system for a telecommunications network, said
telecommunications network including a local switch and a set of terminating
circuits extending between said local switch and terminal equipment provided for
users of the network, each of said terminating circuits passing through a series of
nodes between said local switch and its respective terminal equipment, said fault
25 management system comprising:

circuit testing apparatus located at said local switch and arranged to
perform circuit tests on said terminating circuits; and

a computer system for controlling the circuit testing apparatus, said
computer system including a store containing data relating to said terminating
30 circuits and said nodes;

said method comprising the following steps performed by said computer
system:

instructing said circuit testing apparatus to perform a set of tests on one of said terminating circuits;

checking the results of the set of tests performed by said circuit testing apparatus for the presence of a suspected fault and producing a fault report when
5 a suspected fault is present;

when a fault is suspected, identifying the nodes of the terminating circuit on which the fault is suspected;

evaluating a score for each of said nodes of the terminating circuit on which a fault is suspected which represents the likelihood of the suspected fault
10 being present at the node based on the fault report relating to the suspected fault and using data contained in said data store; and

ranking the nodes on the terminating line on which the fault is suspected in accordance with the scores found in said step of evaluating scores.

This invention will now be described in more detail, by way of example,
15 with reference to the drawings in which:

Figure 1 is a block diagram of an access network and an associated local switch which form part of a telecommunications network in which the present invention may be used;

Figure 2 is a block diagram showing the arrangement of some of the local
20 switches and main switches of the telecommunications network mentioned with reference to Figure 1;

Figure 3 is a block diagram showing the components of the telecommunications network which are used to provide a fault management system embodying this invention for the access network of Figure 1;

25 Figure 4 is a block diagram of the main hardware components of a typical computer;

Figure 5 is a flow chart of the steps which are performed in the fault management system to obtain reference data on a node of the access network;

Figure 6 is a flow chart of the steps which are performed in the fault
30 management system in estimating the location of a suspected fault in the access network;

Figure 7 is a functional block diagram of the fault management system;

Figure 8 is a circuit diagram illustrating some of the measurements which are made when testing a terminating circuit;

Figure 9 is a flow chart of the steps which are performed in the fault management system to monitor the operational condition of the nodes of the
5 access network;

Figure 10 is a graph illustrating how a resistance value is converted into a converted value when performing one of the steps shown in the flow chart of Figure 9; and

Figure 11 is a graph showing some experimental results obtained by
10 monitoring nodes by using the steps shown in the flow chart of Figure 9.

Referring now to Figure 1, there is shown a local switch 10 and a conventional access network 12 connected to the local switch 10. The local switch 10 and the access network 12 form part of a public telecommunications network. The local switch 10 is connected to the terminating circuits or lines of
15 the access network 12. Typically, a local switch is connected to several thousand terminating circuits. Each terminating circuit or line passes through several nodes before reaching its respective terminal equipment. These nodes comprise primary cross-connect points, secondary cross-connect points, distribution points and junctions and examples of these nodes will be described below.

20 In the conventional access network 12 shown in Figure 1, each terminating circuit or line is formed from a pair of copper wires. The copper wires leave the local switch 10 in the form of one or more cables. One of these cables is shown in Figure 1 and indicated by reference numeral 14. The far end of cable 14 from switch 10 is connected to a primary cross-connect point 16 which may be
25 housed in a street cabinet or underground junction box. From the primary cross-connect point 16, the terminating lines branch out as cables in several directions. For simplicity, in Figure 1 there are shown only three cables 18, 20 and 22. The far end of cable 18 is connected to a joint 19. The joint 19 is connected by cable 21 to a secondary cross-connect point 24. The far ends of cables 20 and 22 are
30 connected, respectively, to secondary cross-connect points 26 and 28. For reasons of simplicity, the continuations of the terminating lines beyond secondary cross-connect points 24 and 26 are not shown. The secondary cross-connect

points 24, 26 and 28 are housed in junction boxes which may be located above or below ground.

From the secondary cross-connect point 28, the terminating lines branch out again in several directions in the form of cables. By way of illustration, Figure 1 shows cables 40, 42, and 44 leaving the secondary cross-connect point 28. Cables 40 and 44 are connected, respectively, to joints 46 and 48. Joints 46 and 48 are connected, respectively, to cables 50 and 52, the far ends of which are connected to distribution points 54 and 56.

The far end of cable 42 is connected to a joint 60. The joint 60 is connected by cable 62 to a distribution point 64. For reasons of simplicity, the terminating lines beyond distribution points 54 and 56 are not shown.

Distribution points are implemented as junctions boxes which are typically located on telephone poles. From each distribution point, the terminating lines branch out as single copper wire pairs to where terminal equipment provided for a user of the network is located. By way of illustration, Figure 1 shows two single copper wire pairs 70, 72, leaving the distribution point 64. The far ends of copper wire pairs 70 and 72 are connected, respectively, to terminal equipment 74, 76. As is well known, terminal equipment may take various forms. For example, terminal equipment may be a public telephone located in a public telephone box, a telephone instrument located in a domestic house or an office, or a fax machine or a computer located in a customer's premises.

In the example shown in Figure 1, each of the joints 19, 46, 48 and 60 is used to connect two cables together. Joints may also be used to connect two or more smaller cables to a larger cable.

The cable 14 is housed in a duct. The air in the cable 14 is kept at a pressure above ambient pressure. This discourages ingress of water to the cable 14.

In each terminating line, the two wires of each pair are designated as the A wire and the B wire. At the local switch 10, in order to supply current to the line, a bias voltage of 50V is applied between the A wire and the B wire. As the bias voltage was applied in the early exchanges by using a battery, the bias voltage is still known as the battery voltage. In the terminal equipment, the A wire

and B wire are connected by a capacitor, the presence of which may be detected when the terminal equipment is not in use.

The terminating lines in the access network 10 are prone to faults. The main causes of these faults are ingress of water and physical damage to the nodes through which the terminating lines pass between the local switch 10 and terminal equipment. There are five main faults which occur due to causes arising in the nodes. These faults are disconnection, short circuit, faulty battery voltage, earthing fault and low insulation resistance. A disconnection arises where a terminating line is interrupted between the local switch and the terminal equipment. A short circuit arises where the A wire and B wire of a line are connected together. A faulty battery voltage arises where the A wire or the B wire of a terminating line has a short circuit connection to one of the wires of another line. An earthing fault arises when the A wire or B wire is connected to earth. Low insulation resistance arises where the resistance between the A wire and the B wire or between one of the wires and earth or between one of the wires and a wire of another line is below an acceptable value.

In order to detect faults in the terminating lines of the access network 12, the local switch 10 is provided with a line tester 80. The line tester 80 may be operated from the local switch 10 or, as will be explained in more detail below, from a remote location. The line tester 80 is capable of performing various tests, examples of which will be described below. Various models of line testers for local switches are available commercially. In the present example, the line tester 80 is supplied by Porta Systems of Coventry, England.

Referring now to Figure 2, there are shown some of the switches of the telecommunications network in which the local switch 10 is located. In addition to the local switch 10, Figure 2 shows two main switches 90, 91 and a further local switch 92. The main switches 90, 91 are part of a fully interconnected network of main switches. The local switches 10 and 92 are part of a much larger number of local switches. Each local switch is connected to one or two main switches. Thus, the main switches 90, 91 connect the local switches together.

Referring now to Figure 3, there is shown the local switch 10 and the components of the telecommunications network which provide a fault management system for the access network 12. These components comprise the

line tester 80, a customer service system 100 for the telecommunications network and an access network management system 102.

As shown in Figure 3, the line tester 80 comprises a test head 104 which contains the electronic equipment for physically making line tests and a controller 106 for the test head 104. The controller 106 takes the form of a computer. The controller 106 can be operated from a workstation 108 connected to it and provided at the local exchange 10. The controller 106 is also connected to both the customer service system 100 and the access network management system 102 and can be operated by workstations connected to either the customer service system 100 or the access network management system 102.

The customer service system 100 is also a computer and it can be operated from any one of a number of workstations which are connected to it. In Figure 3, one such workstation is shown and indicated by reference numeral 110. The customer service system 100 is used by operators of the public telecommunications network who have direct telephone contact with the customers of the network. Together with these operators, the customer service system is responsible for providing various services to the customers. These services include the provision of new telephone lines, answering billing enquiries and responding to fault reports received from customers.

The access network management system 102 is also a computer and it can be operated from one of a number of workstations. One of these workstations is shown in Figure 3 and indicated by reference numeral 112. The access network management system 102 is responsible for managing the access network 12 as well as a number of other access networks in the same general geographical area as the access network 12. The access network management system manages various operations for each of the access networks which it manages. These operations include the provision of new equipment, logging data on work performed by engineers in the network, maintaining data on the terminating lines and nodes of each access network detection and management of faults. The workstations which are connected to the access network management system 102 are also connected to the customer service system 100. As shown in Figure 3, the customer service system 100 and the access network management system 102 are connected together.

The operations performed by the customer service system 100 and the access network management system 102 apart from the detection and management of faults in the access network 12 do not form part of the present invention and will not be described in further detail.

5 Although in the present example the fault management system for the access network 12 is formed from the line tester 80, the customer service system 100 and the access network management system 102, the fault management system could also be provided simply by the line tester 80 on its own. In order to achieve this, it would be necessary to add appropriate software to the computer
10 which forms the controller 106. In a small network, this might be an appropriate form of providing the fault management system. However, in a large public network it is advantageous to integrate the fault management system into the customer service system 100 and the access network management system 102.

As mentioned above, the controller 106, the customer service system 100
15 and the access network management system 102 are each implemented as a computer. The main hardware components of a computer are shown in Figure 4. these comprise a central processing unit (CPU) 120, a store 122, a keyboard 124, a visual display device (VDU) 126 and input and output ports 128 for connecting the computer to other computers. The store 122 comprises hard disk storage,
20 floppy disk storage, random access memory and read only memory. The store 122 is used for storing data used by the computer and also for storing the programs which control the computer.

As will be explained below, in the fault management system, one of the parameters which is used for finding a fault in a terminating line, as far as it is
25 available, is the distance of each node through which the terminating line passes from the local switch 10. In a typical access network these distances will be known for some, but not all, of the nodes. For some of these nodes, the distance is known from maps of the cables of the access network. In the present example, where these distances are known from maps, the appropriate data is held both in
30 the customer service system 100 and the access network management system 102.

The fault management system of this example also provides a method of measuring the distance of a node from the local switch 10 when the node is

opened by an engineer. The flow chart for the operations which are performed when measuring this distance is shown in Figure 5 and these operations will now be described.

Initially, in a step 140, an engineer opens a node and makes a telephone
5 call to a colleague sitting at one of the workstations connected to the access network management system 102 to ask him to make line tests. Then, he disconnects one of the line pairs of the node and gives the directory number of this line pair to his colleague in a step 142. Using the connection between his workstation and the line tester 80, the colleague instructs the line tester 80 to
10 perform line tests on the open line pair. In these line tests, the resistance between the A wire and B wire, the capacitance between the A wire and B wire, the capacitance between the A wire and earth and the capacitance between the B wire and earth are measured in a step 144.

Then, in a step 146, the line tester 80 uses the results of the tests to
15 calculate the distance from the switch to the node. Next, in a step 148, the results of the tests including the distance to the node are stored in the customer service system unless the node is a joint. Finally, in a step 150, the results of the tests including the distance to the node are stored in the access network management system 150. These results are stored in step 150 for all nodes
20 including joints.

By performing the operations shown in the flow chart of Figure 5 each time a node is opened, reference data including the distances of the nodes from the switch can be gradually built up. The series of operations shown in Figure 5 is performed under the control of computer programs located in the controller 106,
25 customer service system 100 and access network management system 102.

The controller 106 is programmed to cause the test head 104 to make a series of routine tests each night on each terminating line of the access network 12. These tests will be explained with reference to the circuit diagram shown in Figure 8.

30 In order to test a line, it is disconnected from the switch 10 and connected to the test head 104. Figure 8 shows a line 400 being tested. The line 400 has an A wire 402 and a B wire 404. The end of line 400 remote from switch 10 is connected to terminal equipment 406. Each of the lines 402, 404 has a

resistance which depends upon its diameter and the distance from the local switch to the terminal equipment 406. Each of the wires 402, 404 is coated with an insulating material, which may be a plastic material or paper. The function of the insulating material is to provide insulation between each wire and adjacent wires.

- 5 Damage to the insulating material or oxidation of the metal of a wire can cause the resistance between two adjacent wires to fall.

The effectiveness of the insulation between wires 402, 404 can be determined by measuring the resistance R1 between the A wire 402 and the B wire 404 and the resistance R2 between the B wire 404 and the A wire 402. The
10 resistances R1 and R2 may be different because of rectification as indicated by diodes D1 and D2. For a circuit in good condition, the resistances R1 and R2 are high, greater than 1 megohm. Damage to the insulating material or oxidation will cause the resistances R1, R2 to fall by an amount which depends upon the severity of the damage or oxidation. If the insulating material is totally destroyed
15 so that the A and B wires are physically touching each other, the values of resistances R1, R2 will depend upon the distance between the test head 80 and the point of damage but will typically lie in the range 0 to 1500 ohms. Oxidation can result in wires effectively touching each other.

Only the A and B wires 402, 404 of the line 400 being tested are
20 disconnected. In the other lines, the bias voltage of 50 volts is applied between the A wire and the B wire. In Figure 8, the A wires of the other lines are collectively shown by a wire 410 which is connected at the switch 10 to earth. The B wires of the other lines are collectively shown by a wire 412 connected at the switch to a potential of -50 volts.

25 If the insulating material separating the A wire 402 or the B wire 404 from one of the adjacent A or B wires becomes damaged, or if one of the wires suffers oxidation, current may flow. The effectiveness of the insulation between the A and B wires 402, 404 and adjacent A and B wires can be determined by measuring the resistance R3 between A wire 402 and adjacent A wires 410, the resistance
30 R4 between the A wire 402 and adjacent B wires 412, the resistance R5 between the B wire 404 and adjacent A wires 410, and the resistance R6 between the B wires 404 and adjacent B wires 412.

For a good circuit, the resistance R3, R4, R5, R6 are high, greater than 1 megohm. Damage to insulating material may cause one or more of the resistances R3, R4, R5, R6 to fall by an amount which depends upon the severity of the damage. If the insulating material between the A wire 402 or the B wire 404 and
5 an adjacent wire is totally destroyed so that the two wires are physically touching each other, the resistance between the two touching wires will depend upon the distance between the test head 80 and the point of damage but will typically lie in the range 0 to 1500 ohms. Oxidation can also result in two wires effectively touching each other.

10 The A and B wires 402, 404 and the insulating material between them act as a capacitor. In Figure 8, the capacitance between the A and B wires is shown as having a value C1. The value of the capacitance between the A and B wires of a line will depend upon the length of the line. A break in the line 400 will reduce the value of capacitance C1 as measured from the test head 80. Figure 8 also
15 shows the capacitance C2 between the A wire 402 and earth and the capacitance C3 between the B wire 404 and earth.

Each night, the controller 106 causes the test head 80 to measure the resistances R1, R2, R3, R4, R5, R6 and the capacitances C1, C2, C3 for each terminating line of the access network 12. The controller 106 also causes the test
20 head 80 to check if there is terminal equipment connected to the end of the line. Terminal equipment has a standard capacitance value. When terminal equipment is connected, the value of its capacitance is subtracted from the capacitance as measured by the test head to obtain the capacitance C1. For each terminating line, the results of the tests are stored against its directory number in the access
25 network management system 102.

The controller 106 transmits the results of the tests to the access network management system 102. The access network management system 102 examines the results of the series of tests for each terminating line for the presence of a suspected fault. The possible faults include disconnection, short circuit, a fault
30 battery voltage, an earth fault and low insulation resistance.

When a fault is suspected, the name of the fault and the results of the test for the line are stored in the access network management system 102 against its directory number or an identifier in the exchange associated with the line. The

details of the suspected faults found each night may be reviewed by an operator of the access network management system 102. Where appropriate, the operator may give instructions for a fault to be repaired.

As mentioned above, data relating to the terminating lines and nodes is
5 stored in the access network management system 102. For each node, this data includes the history of engineering interventions in the node, the presence or absence of an active engineering intervention in the node and the presence or absence of an active adverse visual report on the condition of the node. Data on loss of pressure in the main cable 14 is also stored in the access network
10 management system 102.

A mapping of the access network 12 is stored in both the customer service system 100 and the access network management system 102. For each line, and hence for each circuit, the mapping records the route and hence each node through which the line passes. For each node, it records the lines passing
15 through the node.

Referring now to Figure 6, there is shown a flow chart of the series of operations which may be performed when a customer reports a fault on one of the terminating lines of the access network 12 to an operator of the customer service system 100. As will become apparent, in an appropriate case a list is produced of
20 the nodes through which the terminating line passes ranked in order of the likelihood of the fault being present at each node.

Initially a customer reports a fault on one of the terminating lines to an operator in the customer service system in a step 200. If the operator considers that the fault is in the access network 12, the operator sends an instruction to the
25 controller 106 of line tester 80 to perform a series of tests on the terminating line. The line tester 80 then performs the same series of tests as it performs in the nightly tests on the terminating lines. The results of the series of tests are then transmitted by the line tester 80 to the access network management system 102. The access network management system 102 examines the results for the
30 presence of a suspected fault and, where appropriate, estimates the distance of the fault from the local switch 10. The details of the suspected fault are transmitted to the customer service system 100.

Where the operator in the customer service system considers that it is appropriate for the fault to be investigated by an operator in the access network management system 102, in a step 201, the operator in the customer service system asks an operator in the access network management system 102 to
5 investigate the fault.

The operator in the access management system then decides what to do about the fault report. This decision will depend on the details of the suspected fault and his knowledge of the state of the terminating lines in the access network 12. For example, if the fault report indicates that there is a disconnection and the
10 operator knows that one of the nodes through which the terminating line passes has been severely damaged but is under repair, it may be appropriate for the operator merely to advise the customer that the damage is being repaired. Usually, the operator will decide to obtain a list of the nodes through which the terminating line passes ranked in the order of the likelihood of the fault being present at each
15 node. If the operator decides to obtain this list, in a step 203, the operator invokes a ranking algorithm. The remainder of the flow chart shown in Figure 6 refers to the steps which are performed in this algorithm.

The ranking algorithm is entered in a step 204. The ranking algorithm is held as part of a program of the access network management system 102. As will
20 become apparent from the following description, for each of the nodes on the terminating line on which a fault is suspected, an individual score is evaluated for each one of a set of factors relating to the likelihood of the suspected fault being present at the node and the individual scores are then combined to produce a combined score for the node. The nodes are then ranked in accordance with their
25 combined scores. In evaluating each individual score, the data contained in the fault report and also the data contained in the access network management system relating to the nodes and the terminating lines are used.

After entering the ranking algorithm, in a step 205, the route followed by the terminating line on which a fault is suspected is identified and all the nodes on
30 this route are then found. Next, in a step 206, the first node on this route is selected.

In a step 207, the individual score for the first factor is evaluated. The first factor is historical intervention analysis. The score for this factor is evaluated in five stages as follows:

stage 1: for at least one engineering intervention in previous five days,
5 score 3000;

stage 2: for at least one engineering intervention in the previous 30-day period but excluding the previous five-day period, score 2000;

stage 3: for at least one engineering intervention in the previous 90-day period but excluding the previous 30-day period, score 1000;

10 stage 4: for the total number of engineering interventions in the previous 90-day period, score 100 times number of interventions.

The total score for the historical intervention analysis factor evaluated in step 207 is then the sum of the scores found in stages 1 to 4 above.

The next factor is visual inspection analysis and this is evaluated in step
15 208. If there is an adverse active visual report for the node, the score for this factor is evaluated in step 208 as 2000.

The next factor is active intervention analysis and this factor is evaluated in step 209. This factor is evaluated in two stages. In the first stage, if there is active engineering work at the node, a score of 3000 is made. In the second
20 stage, if there is active engineering work on any terminating line passing through the node, a score of 3000 is made. The total score for this factor is then the higher of the scores found in stages 1 and 2.

The next factor is suspect line analysis and this factor is evaluated in step 210. For each terminating line which passes through the node and on which a
25 fault is suspected and which is of the same type as the suspected fault under study, a score of 1000 is made. Thus, if the suspected fault is a disconnection and disconnection is also suspected on three other terminating lines which pass through the node, then the total score is 3000.

Next, in a step 211, a check is made for the presence of scoring factors
30 among the visual inspection analysis factor as evaluated in step 208, the active intervention analysis factor as evaluated in step 209 and the suspect line analysis factor as evaluated in step 210. If two or three of these factors score, then the combined score for the factors evaluated in steps 208 to 210 is the sum of the

individual factors multiplied by two. If only one of these factors scores, then the total score for the factors evaluated in steps 208 to 210 is the score for this individual scoring factor.

Next in a step 212, the score is evaluated for the proximity analysis factor. The score for this factor is evaluated in two stages. As will be recalled from the above description, the access network management system contains data on the distance between the local switch 10 and the nodes derived from two sources. The first source is the measurements which are made as a result of engineers opening the nodes and performing the series of operations shown in the flow chart of Figure 5. The second source of information is the distances which are found from maps of the terminating lines. The proximity analysis is performed using both sources of data.

For the first source of data, that is distances derived from opening the nodes, the score is evaluated as follows:

if the proximity of the node to the estimated location of the suspected fault is less than 50 m, score 3000;

if the proximity of the node to the estimated location of the suspected fault is equal to or greater than 50 m but less than 100 m, score 2000,

if the proximity of the node to the estimated location of the fault is equal to or greater than 100 m and less than 150 m, score 1000.

For the second source of data, that is data derived from maps of the terminating lines, the score is evaluated in the same manner as that specified for the first source of data. The total score for the proximity analysis factor is then obtained by adding the scores obtained by using the data from the two sources.

In step 212, when evaluating the score for each source of data, if the distance is not available for the distance of the node from the switch for that source of data, then the score is 0.

Next, in a step 213, a score is evaluated for the pressure analysis factor. The pressure is the pressure in the cable 14 leading from the switch 10 to the primary cross-connect point 16. If the pressure is below a threshold value, the score is evaluated as 3000.

Next, in a step 214, the operator is invited to add a manual score and, normally, an operator will do this if he knows of some special problem.

Next, in a step 215, each scoring factor is compared with the fault type and an additional score is made in accordance with Table 1 set out below. In this table, the visual inspection analysis factor of step 208, the active intervention analysis factor of step 209 and the suspect line analysis factor of 210 are treated
5 as a single factor. The proximity analysis factor evaluated in step 212 is divided into two sub-factors. The first sub-factor is the proximity analysis performed using the first source of data, that is data obtained by opening the node and calculating the distance from line tests. The second sub-factor is the proximity analysis factor evaluated from the second source of data, that is maps of the terminating lines.

10 Referring to Table 1, for example, if the fault is a short circuit, additional scores are made as follows. If there is a score for the historical intervention factor of step 207, an additional score of 3000 is made. If there is a score for one or more of the factors of steps 208 to 210, an additional score of 1000 is made. If there is a score for the pressure analysis factor of step 213, an additional score of
15 4000 is made. If there is a score for the proximity analysis factor derived from the first source of data, a score of 4000 is made. If there is a score for the proximity analysis factor obtained from the second source of data, a score of 4000 is made. The total additional weighting score is the sum of the individual scores.

TABLE 1
Additional Weighting for Fault Type

	Disconnection	Short Circuit	Battery	Earth	Low Insulation Resistance
Historical intervention	3000	3000	3000	3000	3000
Visual inspection	1000	1000	4000	4000	4000
Active intervention					
Suspect line					
Pressure	4000	4000	2000	2000	2000
Proximity (first data source)	4000	4000	2000	2000	2000
Proximity (second data source)	4000	4000	2000	2000	2000

5 Next, in a step 216, a check is made for coincidence in scoring factors. More specifically, if there is only a single scoring factor, a jump is made to step 218. If there are two or more scoring factors, the algorithm goes to step 217 before proceeding with step 218. For the purposes of step 216, the factors of steps 208 to 210 are combined and the factor of step 212 is divided into sub-
10 factors in the manner discussed in relation to Table 1 above.

 In step 217, if there are two or more scoring factors, the highest score is taken and increased by 10000. For this purpose, for steps 208 to 211, the score obtained for these steps at the end of 211 is used. The score obtained in this way then becomes the combined score for the individual factors for the node. If there
15 is only a single scoring factor, this factor becomes the combined score for the nodes.

 In step 218, a check is made to determine if there are any more nodes for which a combined score is to be evaluated. If there are one or more nodes still remaining, the algorithm returns to step 206 and the next node is selected. If
20 there are no more nodes, the algorithm continues with step 219.

 In step 219, the nodes are ranked in order of their scores in descending order. Thus, the node at the top of the list is the most likely location of the fault. In step 210, the node with the highest score is then selected and in step 211 the

operator is given the opportunity to override the result and an operator would normally do this in view of special circumstances known to him.

Referring now to Figure 7, there is shown a functional block diagram of the fault management system described with reference to Figures 3 to 6. As shown in Figure 7, the fault management system comprises the test head 104, the test head controller 106, a test report analyser 300, a store 301 and a node identifier 302. The test head 104 and the test head controller 106 have already been described. The test report analyser 300, the store 301 and the node identifier 302 are implemented by the access network management system 102.

10 More specifically, the test report analyser is the part of the program of the access network management system which is responsible for performing steps 206 to 221 of the ranking algorithm. The store 301 is that part of the store of the computer used to implement the access network management system 102 and which contains data relating to the nodes and the terminating lines. The node

15 identifier 302 is that part of the program which is responsible for step 205 of the ranking algorithm.

There will now be described with reference to Figure 9 the individual steps of a routine which is performed each night by the access network management system 102 to monitor the operational condition of the individual

20 nodes of the access network 12. As will become apparent from the following description, the routine uses the resistance values R1 to R6 for each circuit obtained by the routine nightly test on the individual circuits. The routine shown in Figure 9 also uses the mapping of the access network 12 which is stored in the access network management system 102.

25 Referring now to Figure 9, after entering the routine, in a step 500 a circuit score S is calculated for each circuit from the resistance measurements R1 to R6. The circuit score S is indicative of both the probability of a circuit having or developing a fault and the operational quality of the circuit. A relatively high circuit score indicates that there is a high probability of the circuit having or

30 developing a fault and that the operational quality of the circuit is poor. A relatively low circuit score indicates that there is a low probability of the circuit having or developing a fault and that the operational quality of the circuit is good.

In order to calculate the circuit score S for a circuit, each resistance value R_i is converted into a converted value V_i using a formula which will be described below. The converted value V_i is indicative of the probability of the resistance causing a fault. Each converted value V_i is multiplied by a weighting factor W_i to obtain the product $V_i * W_i$. The products $V_i * W_i$ for all six resistance measurements are then summed to produce the circuit score S.

Thus, the circuit score of S is defined by the following equation:

$$S = \sum V_i * W_i$$

10

Each converted value V_i is calculated for the corresponding resistance value by using the following formula:

If $R_i \leq P$, $V_i = 0$
 If $P < R_i \leq Q$, $V_i = 1$
 If $R_i > M$, $V_i = 0$
 If $Q < R_i \leq M$, then

20

$$V_i = \{ [1 / (1 + ((R_i - Q) / L)^{0.3})] + [-1 * (R_i - 1 * 10^6) / (1 * 10^6 - Q)] \} / 2$$

where P is a lower threshold set to 1 k Ω
 Q is an upper threshold set to 5 k Ω
 M = 1 M Ω
 L is constant set to 1 k Ω

25

A graph of V_i plotted against R_i using the formula set out above is shown in Figure 10.

The formula for calculating V_i has been developed empirically. However, the physical explanation behind the formula is as follows:

When a resistance R_i of a circuit has a value less than 1 k Ω , it is very likely that the cause of the low resistance is a fault in the terminal equipment, and so the value of V_i is set to zero. Where a resistance R_i has a value in the range

1k Ω to 5k Ω it is almost certain that there is an associated fault. For a resistance value greater than 5k Ω , the probability of the existence of an associated fault falls both progressively and progressively less steeply with increasing value of R_i . Where a resistance R_i has a value greater than 1M Ω , it is very unlikely that there is an associated fault.

For each type of resistance R_1 to R_6 , the value of the corresponding weighting factor is determined from the following table:

<u>Types of Resistance</u>	<u>Weighting Factor R_i</u>
R1 (A wire to B wire)	15
R2 (B wire to A wire)	15
R3 (A wire to earth)	5
R4 (A wire to -50V)	30
R5 (B wire to earth)	30
R6 (B wire to -50V)	5

After determining the circuit score S for each circuit in step 501, the routine progresses to step 502. In this step, it calculates a node score H for each node.

In order to calculate the node score H for a node, all the circuits passing through the node are identified. The individual circuit scores S for the individual circuits passing through the node are then summed and the result is divided by $\sqrt{n}$. n is the number of resistance measurements for circuits passing through the node having a value less than 1M Ω . Thus, the node score H for a node is defined by the following equation:

$$H = \sum S \cdot \sqrt{n} / n = \sum S / \sqrt{n}$$

The node score H for a node is indicative of both the probability of the operational state of the node causing a fault in one or more of the circuits passing through the node and the operational quality of the node. A relatively high node score indicates that there is high probability of the operational state of the node causing a circuit fault and that the operational quality of the node is poor. A relatively low node score indicates that there is a low probability of the operational

state of the node causing a fault and that the operational quality of the node is good.

In the equation set out above for calculating a node score H , ΣS is divided by n so as to provide an averaging effect for resistance measurements on circuits passing through the node having a value less than $1\text{M}\Omega$. This makes node scores for nodes carrying a large number of circuits comparable with node score for nodes carrying a small number of circuits. As will be explained below, the nodes are ranked in accordance with their node scores. However, the node scores of nodes carrying only a few circuits are very sensitive to the number of resistance measurements less than $1\text{M}\Omega$. More generally, as the number of circuits passing through a node increases, and hence potentially the number of resistance measurements less than $1\text{M}\Omega$ increases, the reliability of the node score as an indicator of the operational quality of the node increases. In the equation set out above for calculating a node score H , $\Sigma S/n$ is multiplied by $\sqrt{n}$ so as to give progressive emphasis to the node score as the number of resistance measurements less than $1\text{M}\Omega$ increases.

Next, in a step 502, the nodes are ranked in accordance with their node scores H . The node having the highest node score H is identified as the worst node and selected for further investigation.

Although the relatively high node score of the worst node may be caused by the operational state of the worst node, it may also be caused by the operational state of one of the nodes on the route from a local switch to the worst node. The node score H is obtained from resistance measurements and does not take other factors into account. In order to identify the exact location of the node having a poor operational quality, the ranking algorithm described with reference to Figure 6 is used.

Accordingly, in step 503, a route from the local switch 10 to the worst node is found and all the nodes on the route up to, and including, the worst node are found. Steps 206 to 219 of the ranking algorithms are then executed for this route in a step 504.

Occasionally, there is more than one route from a local switch to a node. In a step 505, a check is made for the existence of another route to the worst

node. If there are one or more further routes, steps 505 and 504 are executed for all such further routes.

Lastly, in a step 506, the operator decides whether to ask an engineer to investigate on or more of the nodes for possible faults. When the routine reaches
5 step 506, the operator will know the identity of the worst node. Also, for the nodes on each route to the worst node, he will have a list of the nodes as ranked by the ranking algorithm. In deciding whether to ask an engineer to investigate one or more of the nodes, the operator combines this data with his own knowledge of the access network 12.

10 In addition to investigating the worst node, the operator may also investigate other nodes, for example the next worst node or other nodes having high node scores. In order to do this steps 504 to 506 are performed for each such node.

The routine of Figure 9, make it possible to identify and investigate at an
15 early stage nodes where there is a risk that the operational state of the node is deteriorating to a point where faults may occur. It will usually be possible to restore the operational state of such node before a fault occurs which leads to a fault report.

Figure 11 shows the results of some experimental work in the nodes in an
20 access network forming part of BT's UK public telecommunications network. The node score was evaluated for a large number of nodes. Each node was then monitored for a fault report from a customer during the subsequent 3 months. In Figure 11, for these nodes, the node score is plotted against fault reports received during the three months after evaluating the node score. These experimental
25 results show a strong correlation believed node score and number of fault reports.

Although the present invention has been described with reference to an access network in which each circuit is carried by a piece of copper wire, it may also be used for terminating circuits carried by optical fibres..

**CLAIMS**

1. A fault management system for a telecommunications network including a local switch and a set of terminating circuits extending between said local switch and terminal equipment provided for users of the network, each of said terminating circuits passing through a series of nodes between said local switch and its respective terminal equipment, said fault management system comprising:
- 5 circuit testing apparatus located at said local switch and arranged to perform circuit tests on said terminating circuits;
- 10 a store containing data relating to said terminating circuits and said nodes;
- means for instructing said circuit testing apparatus to perform a set of tests on one of said terminating circuits;
- means for checking the results of a set of tests performed by said circuit testing apparatus for the presence of a suspected fault, said checking means being
- 15 arranged to produce a fault report when a fault is suspected;
- means for identifying the nodes of a terminating circuit on which a suspected fault is present;
- means for evaluating a score for each node on a terminating circuit line on which a fault is suspected which represents the likelihood of the suspected fault
- 20 being present at the node, said evaluating means using the fault report relating to the suspected fault and data contained in said data store in evaluating the score;
- and
- means for ranking nodes of a terminating circuit on which a fault is suspected in accordance with their scores found by said evaluating means;
- 25 whereby, in use, after testing a circuit on which a fault is suspected, a list is produced of the nodes on said circuit in which the nodes are ranked in accordance with the likelihood of the fault being present at each node.
2. A fault management system as claimed in claim 1, in which, for each node
- 30 on a terminating circuit on which a fault is suspected, said evaluating means is arranged to evaluate an individual score for each one of a set of factors and to combine said individual scores to produce a combined score representative of the likelihood of the suspected fault being present at the node.



3. A fault management system as claimed in claim 2, in which said set of factors includes the history of engineering interventions at the node.
- 5 4. A fault management system as claimed in claim 2 or claim 3, in which said set of factors includes the presence or absence of active engineering intervention at the node.
5. A fault management system as claimed in any one of claims 2 to 4, in
10 which said set of factors includes the presence or absence of a current adverse visual report on the node.
6. A fault management system as claimed in any one of claims 2 to 5, in which said set of factors includes the presence or absence of a suspected fault on
15 another terminating circuit passing through the node.
7. A fault management system as claimed in any one of claims 2 to 6, in which said fault report includes an estimate of the location of the fault and said set of factors includes the proximity of the estimated location of the fault to the
20 location of the node.
8. A fault management system for a telecommunications network, said telecommunications network including a local switch and a set of terminating circuits extending between said local switch and terminal equipment provided for
25 users of the network, each of said terminating circuits passing through a series of nodes between said local switch and its respective terminal equipment, said fault management system comprising:
- 30 circuit testing apparatus located at said local switch and arranged to perform circuit tests on said terminating circuits;
- a computer system for operating said circuit testing apparatus;
- said computer system including a store containing data relating to said terminating circuits and said nodes;



said computer system being controlled by at least one program so as to perform the following operations:

instruct said circuit testing apparatus to perform a set of tests on one of said terminating circuits;

5 check the results of a set of tests performed by said circuit testing apparatus for the presence of a suspected fault and produce a fault report when a suspected fault is present;

identify the nodes of a terminating circuit on which a suspected fault is present;

10 evaluate a score for each node on a terminating circuit on which a fault is suspected which represents the likelihood of the suspected fault being present at the node by using the fault report relating to the suspected fault and data contained in said store; and

rank nodes of a terminating circuit on which a fault is suspected in
15 accordance with their scores.

9. A method of operating a fault management system for a telecommunications network, said telecommunications network including a local switch and a set of terminating circuits extending between said local switch and
20 terminal equipment provided for users of the network, each of said terminating circuits passing through a series of nodes between said local switch and its respective terminal equipment, said fault management system comprising:

circuit testing apparatus located at said local switch and arranged to perform circuit tests on said terminating circuits; and

25 a computer system for controlling the circuit testing apparatus, said computer system including a store containing data relating to said terminating circuits and said nodes;

said method comprising the following steps performed by said computer system:

30 instructing said circuit testing apparatus to perform a set of tests on one of said terminating circuits;



checking the results of the set of tests performed by said circuit testing apparatus for the presence of a suspected fault and producing a fault report when a suspected fault is present;

when a fault is suspected, identifying the nodes of the terminating circuit
5 on which the fault is suspected;

evaluating a score for each of said nodes of the terminating circuit on which a fault is suspected which represents the likelihood of the suspected fault being present at the node based on the fault report relating to the suspected fault and using data contained in said data store; and

10 ranking the nodes on the terminating line on which the fault is suspected in accordance with the scores found in said step of evaluating scores.

10. A method of operating a fault management system as claimed in claim 9, in which, in said step of evaluating a score for each of said nodes, for each of said
15 nodes an individual score is evaluated for each one of a set of factors and the individual scores are combined to produce a combined score representative of the likelihood of the suspected fault being present at the node.

11. A method of operating a fault management system as claimed in claim 10,
20 in which said set of factors includes the history of engineering interventions of the node.

12. A method of operating a fault management system as claimed in claim 10 or claim 11, in which said set of factors includes the presence or absence of active
25 engineering intervention at the node.

13. A method of operating a fault management system as claimed in any one of claims 10 to 12, in which said set of factors includes the presence or absence of a current adverse visual report on the node.

30 14. A method of operating a fault management system as claimed in any one of claims 10 to 13, in which said set of factors includes the presence or absence of a suspected fault on another terminating circuit passing through the node.



WO 97/19544

PCT/GB96/02756

15. A method of operating a fault management system as claimed in any one of claims 10 to 14, in which the fault report includes an estimate of the location of the fault and said set of factors includes the proximity of the estimated location of
- 5 the fault to the location of the node.