

- 1 -

IN-SERVICE MONITORING IN COMMUNICATION NETWORKS

This invention relates to means for monitoring of digital communication networks, in particular connections within the network which do not have a deterministic frame structure.

It is desirable in most, if not all, communication networks, that there should be some form of in-service monitoring, so that failure of a part of the network is detected promptly and can be dealt with. There are many ways in which a failure may be detected and some form of indicator is normally required. Digital communications systems used by network operators have an established practice in which a loss of signal detected by a piece of equipment in a path and forming part of a network causes the equipment to send a special signal to indicate loss of signal in the downstream direction. This known form of special signal, the Alarm Indication Signal, or AIS, is normally in the form of a continuous stream of digital "1"s. The AIS could be in any other form defined and recognised by the network in question, but digital "1"s have been found to be the most convenient to implement. A key role of the AIS is to prevent multiple alarms being raised in downstream equipment. The AIS detection criteria can be set at a threshold of a given number of continuous digital "1"s, or a given density of "1"s over a certain span of bits. The AIS may be set at a relatively high level so that the alarm is not triggered after only a short string of 1s. Because network performance is important to both customer and network operator, both the time and the number of incidents for which an AIS is triggered may be recorded.

Communication paths are available at various bit rates or rates of data transmission to suit the customer's requirements. At the high data rates there is sometimes a deterministic frame structure. A deterministic frame



structure is one in which a source sends a signal with a well defined structure at the binary level in the form of overhead plus payload. At lower data rates, such as a 64Kbit/sec there tends to be no generic deterministic frame structure.

'Structured' services usually take the form of a standardised path overhead (such as a periodic framing pattern perhaps with a CRC-type function which can be used for monitoring purposes) and a well defined payload area for the, usually non-deterministic, customer traffic signal. An example of a structured signal is the G.704/G.706 ITU recommended frame structure for the 2048kbit/s network layer path signal. In this example the path overhead is defined in Timeslot 0 (which contains, amongst other functions, both a frame alignment signal and a CRC), and the payload the customer can use is, in most cases, the remaining 31 Timeslots of the frame in some format (thus giving the customer a maximum aggregate 1984kbit/s payload channel).

Many network operators offer services which take the form of giving the customer the whole bit rate of the digital path in question. For example, this could be a $n \times 64$ kbit/s path (where $n = 1, 2, 3 \dots$ to some limit perhaps) or a higher bit rate hierarchical network layer, such a 2048kbit/s or 34368kbit/s path. This creates a problem for the network operator offering such a service because there is no generic frame structure requirement imposed on the signal at the binary level (either by the customer or the network operator) which can be used by the network operator to monitor the end-to-end performance of the path at the binary level. In this specification such service offerings will be referred to as 'unstructured'.

A particular problem with unstructured services is that the network operator might not be quickly (if at all) aware of path failures. In some cases this information is first provided by the customer.



In public telecommunications networks systems where quality of service and continuity of service are specified to very high levels, if not guaranteed, an extra path can be provided for back-up purposes. This extra path will generally take a different route from the primary path to assist in providing the level of service being supplied to the customer. More generally only one (traffic carrying) path will be provided on an end to end basis because the provision of an additional end to end path would significantly increase the cost of the standard service.

For reasons of economy, it is therefore desirable that capacity for a standard service in-service monitoring be derived purely from the path which carries the traffic. For instance, as mentioned above, it is known to use an Alarm Indication Signal (AIS), which is commonly a signal comprising all "1"s. This particular all "1"s AIS is widely used (and defined in the international standards bodies such as the ITU) because for one reason it is easily implemented in practice in a wide range of terminal equipment.

According to one aspect of the present invention there is provided a digital communications path having means for monitoring the path, the path having two ends, a first end receiving incoming traffic signals and connected to a scrambler device for scrambling said incoming traffic signals, a transmitter device for transmitting scrambled signals along the path, and a second end at which the transmitted signals are received from the first end, the second end being provided with monitoring means for detecting an Alarm Indication Signal in the transmitted signals and unscrambling means for unscrambling the transmitted signals.

The AIS may, for instance, comprise all "1"s, in a digital communications system. Embodiments of the present invention can then provide communication paths dedicated to the use of a specific customer wherein the customer's signal



is scrambled at the input to the communications path and unscrambled before being received by the customer's equipment at the end of the path. This allows the customer ostensibly to use an all "1"s signal for his own purposes, such a signal in practice being scrambled while carried by the path and therefore not detectable as an AIS by the monitoring means at the receiving end of the path.

According to another aspect of the present invention there is provided a method of monitoring a digital communications path including means for receiving traffic signals to be transmitted along the path at a first end of the path, scrambling the traffic signals by use of a scrambler device, transmitting said scrambled traffic signals along the path, at a second end of the path monitoring the path for an AIS in the scrambled traffic signals, indicating the presence of an AIS in the event of an AIS being detected in the scrambled traffic signals, and unscrambling said received traffic signals by means of an unscrambling device.

In an embodiment of this invention, a customer's data may be scrambled as it transits a NTU (Network Terminating Unit) at a first end of the path and descrambled as it transits a second NTU at a second end of the path. The NTU is used to connect individual customers data to a telecommunications network.

Many of the paths currently operating at any network level, for example $n \times 64\text{kbit/sec}$, 2048kbit/sec , 8448kbit/sec and 34368kbit/sec and provided as a standard service, do not have a deterministic frame structure comprising payload and overhead. Embodiments of this invention could be used to provide performance monitoring functions for these paths.

There are generally two types of network failure: a first is the 'hard' failure, where a true loss of a signal (at some network layer) occurs and for which AIS will persist in all downstream client layers until either automatic or manual



restoration is effected; and a second is the 'soft' failure, where a transient but gross disturbance of a network layer signal occurs and that network layer loses path alignment. In such a case both the affected network path layer and all client path layers go through a complex loss and recovery of frame alignment until normal operation is resumed. The key point to note here is that the network self-recovers without any form of intervention.

The first type of failure is rather obvious and its implications are easily understood. The second type of failure is, however, more interesting and its implications need to be clearly understood to optimise the solution.

The soft failure is usually a consequence of short but intense error bursts in networks (for example, as small as 10µs at 140Mbit/s). A key facet of this effect, is that each client layer loses and recovers frame alignment due to corruption of the justification process, which either inserts an extra bit or deletes a traffic bit from the client layer signal. This effect is known as an 'uncontrolled bit-slip'.

Examining the effect of an uncontrolled bit-slip as it ripples through the successive client network layers shows a complex pattern of random error events (with a bit error density probability, $p_e = 0.5$) intertwined with very short bursts of AIS. The AIS occurs due to the detection of loss of frame alignment at a given layer and persists until such time as the layer in question recovers frame alignment.

It is generally accepted by the man skilled in the art of error performance that the above described soft failure phenomenon can be interpreted as giving rise to a SES - both at the network layer of origin and in all supported client network layers. In the case of a hard failure, AIS is usually a permanent state (until manual or automatic restoration is effected) and hence SES events are easily measured.



One of the most important parameters to measure from the perspective of both the network operator and the customer is Unavailability. The customer and network operator might also have contractual performance agreements on 'Service Availability', such that it is highly desirable that such a measurement can be performed accurately in-service.

A path is usually defined to enter the Unavailable state (from the Available state) at the onset of a period of several consecutive Severely Errored Seconds (SES). Within the ITU Recommendations on error performance, namely G. 821 and G. 826, this is currently defined to be 10 consecutive SES. Exit of the Unavailable state (to the Available state) is usually defined at the onset of a period of several consecutive non-SES. Within the aforementioned ITU Recommendations this is currently defined to be 10 consecutive non-SES.

A further parameter which has been found important to measure is 'Short-Breaks' (SB). A SB is defined in this patent as a number of consecutive SES which terminate before the threshold which defines the onset of Unavailability. A SB event terminates on the first non-SES. For example, and in consideration of the above discussion, a SB could be defined as a period of between say 3 and 9 consecutive SES. Moreover, such events could be measured as a single class, or they could be sub-divided into several class intervals, with the frequency of SB events in each class interval measured.

From the above it is clear that the SES (also defined in the above named ITU Recommendations) is the key primary parameter to measure.

According to a further embodiment of the present invention there is provided a network terminating unit (NTU) for use in a digital communications path having means for receiving an incoming digital traffic signal, scrambler means for scrambling the traffic signal, and transmission means for transmitting the scrambled signal along a path.



In order to implement a network of the type described it is possible to use a pair of NTUs, with one connected to each end of a digital path.

If the availability and error performance is to be recorded a method of measuring it needs to be devised. Since the key measurements are Unavailability and Short Breaks, and since both are based upon the SES, this measure could be the occurrence of the SES. Performance information can either be reported immediately to a network management centre as it occurs, or stored at the NTU where it is measured and then reported once say, every 24 hours, either in response to a polled request or by means of a known autodialling system. Whether immediate or deferred reporting is required will depend upon the nature of the impairment and the maintenance philosophy of the network operator. For example, it could be decided that all unavailability events are always immediately reported (and perhaps stored), whilst occasional SES events are generally stored and only reported in response to polling from a network management system.

The size of any scrambler/descrambler used will affect network performance. If there is a disturbance at some point then the greater the number of stages in the scrambler the greater the time taken for the disturbance to filter through the descrambler and the greater the time to settle back to normal operation after the disturbance. There are therefore distinct advantages in using a scrambler with a small number of stages.

Using a scrambler with a small number of stages does lead to the risk that there is a possibility of the combination of incoming data for a particular state of the scrambler could generate a sequence of digital "1"s of greater length than the threshold length n which is set to detect an AIS. In such circumstances the detection of the AIS would be false. In order to overcome this and prevent such an event occurring it



is proposed to incorporate an Adverse State Detector which can be used to move the scrambler to a different state. The Adverse State Detection system checks the binary sequence input into the network by the NTU downstream of the scrambler. If a persistent stream of digital "1"s is detected then the adverse state detector invokes a change of state of the scrambler. By use of a scrambler with a relatively simple generator polynomial, such as $G(x) = x^2 + x + 1$, the scrambler can be implemented in practice in only four stages. This particular function, $G(x)$, is an example of a primitive polynomial.


The invention will now be described in detail with reference to the attached drawings in which:

Figure 1 shows a general configuration of a single network connection with facilities for in-service performance monitoring, and

Figure 2 shows an example of a self-synchronous scrambler/descrambler with an example of a type of adverse state detection mechanism, and

Figure 3 is a table showing an example of adverse state detection on a scrambler/descrambler,

Figure 1 shows a simple network in which the signal flow is described from left to right. (Similar considerations apply to signals flowing in the reverse direction). A piece of Customer Premises Equipment (CPE) 1 is connected via a path 2 to a Network Terminating Unit (NTU) 3. The signals are then scrambled in a scrambler 4 and transmitted along a path 5 of a serving network 6, before passing out at 7 to a second NTU 8. At a second end of the path NTU 8, includes an AIS detector and measurement unit 9. The detector and measurement unit 9 detects the incoming signals and if an AIS, a continuous sequence of digital "1"s, is detected for longer than a predetermined period, an AIS is deemed to have been detected. Any one second time period which has an AIS



condition present in it is also deemed to be a SES. If no AIS is detected (i.e. there is not a continuous sequence of "1"s which exceed a preset threshold n) then the signal is passed to a descrambler 10 where it is descrambled before passing out of the NTU 8, along 11 to a second CPE 12. A similar process occurs for the signals transmitted in the other direction, i.e. from right to left of the figure. This embodiment enables a customer to use the full bandwidth of the path with a performance monitoring function but without the extra expense of providing additional bandwidth for the performance monitoring.

Reporting performance data and/or specific events to the Network Management System (NMS) 14 may be, for example, by a dial-up PSTN link (public switched telephone network), by dial up ISDN (Integrated Services Digital Network) or by X25. Once the reports have been received by the NMS 14 then appropriate action can be taken.

On unstructured digital paths 5, 6, 7 the customer can generate any binary sequence of '1s' and '0s' usually without restriction. Hence, if the customer signal can be prevented from generating a sequence of greater than say n consecutive binary '1s' into the serving network 6 (i.e. from transmitting NTU 3), and n is chosen optimally in respect of the foregoing discussion, then if a sequence of greater than n consecutive binary all '1s' is detected at the receiving NTU 8 this can be attributed as sourced from within the serving network infrastructure 6 (i.e. a true AIS signal) and can be associated with a SES event occurring.

A scrambler device 4 can be incorporated into the transmitting section of the NTU such that it scrambles the customer data signal before it is sent onward into the serving network 6. The receiving section of the NTU would incorporate a corresponding descrambler device 10 to recover the received network signal at the distant end of the path 7 and convert



it back into the form originally sent by the customer at the transmitting end of the path. This signal would then be delivered to the CPE 12 at the receiving end of the path 11.

The probability of the scrambler generating a consecutive sequence of binary '1s' is a function of both the input customer signal sequence and the state of the scrambler at any point in time. This probability can, perhaps, be made arbitrarily small by using a scrambler (and descrambler) with many stages. However, a large number of stages in the (self-synchronising) scrambler/descrambler equates with a proportionately increased delay in regaining scrambler/descrambler synchronisation (say due to error events) and there could be increased cost implications. Moreover, the probability of the scrambler not generating a sequence of greater than say n binary '1s' cannot be reduced to zero in any case.

One method of ensuring that the scrambler cannot generate a consecutive sequence of greater than n binary '1s' is to use an adverse state detector mechanism within the design of the scrambler 4. This mechanism checks the binary sequence sent into the network 5, 6, 7 from the NTU 3, and if it detects a persistent stream of all '1s' the mechanism invokes a change of state of the scrambler 4. A similar system can be used at the descrambler 10 to correct for this aberration (though this can be omitted on cost grounds if required - the implications of doing this are covered later).

Figure 2 is an example of a suitable scrambler. This scrambler operates a small number of stages. This is possible because a simple generator polynomial has been used. It is of the form $G(x) = x^4 + x + 1$. This simple polynomial can be implemented in just four stages. The small number of stages is advantageous in ensuring the descrambler regains synchronisation in a short period - 4 clock periods for the example given - after the disturbing event ends.



The example adverse state detector shown consists of an AND gate 20, which logically summates the outputs from the three delay stages x , x^2 and x^3 of the scrambler/descrambler. The AND gate is assumed to be at logic state 1 when all three delay stages are at logic state 1. A summation device Σ counts the number of clock cycles that the outputs of the three delay stages stay at logic state 1. An exclusive OR gate, shown as $\oplus$ between delay stages x^2 and x^3 has the function of inverting the value of the delay stage x^2 . So if during operation of the scrambler all four delay stages are set at 1, and the customer data input signal is also 1, then the output signal to the network will also be 1. This prevails for so long as the customer signal remains at 1. The summation device will then count the number of clock cycles that delay stages x , x^2 and x^3 remain at 1. Whilst this remains below a preset threshold the output logic state of the summation device remains at 0. Once the threshold is reached the summation device output is set to logic state 1 which, on the next clock cycle, sets the value in delay stage x^2 to 0. The action of reaching threshold causes the summation device to reset its count and output stages. With complimentary adverse state mechanisms at both scrambler and descrambler, the aberration is corrected at the descrambler. An example of this scrambler and the particular sequence of operation is given in Figure 3.

At cycle $n+4$ it can be seen that the scrambler/descrambler delay stages are all set to '0'. If, at this point, the customer data became a persistent sequence of consecutive '0s' then the NTU output to the network would also become a sequence of '0s' for as long as the customer sequence of consecutive '0s' persisted. This certainly does not present a problem for the invention as claimed, nor should it be a problem for the network, assuming that all network



transmission systems encountered have been designed to cope with this, i.e. there are no spectral or timing issues.

Although not proved here, it can easily be shown that whatever initial state the scrambler and the descrambler are in (e.g. both different), they both become synchronised (i.e. have the same state) under error-free conditions after a maximum of 4 clock cycles in the example given - or more generally, after a number of clock cycles equal to the number of delay stages in the scrambler/descrambler. This is indeed how the self-synchronous scrambler/descrambler recovers after any period of gross disturbance, such as a large error burst or a break in transmission.

It should be noted that any single transmission error gets multiplied by the number of feedback taps of the descrambler. In most cases this will be 3 as in the example given above. Hence, one could argue that $G(x)$ should also be a primitive polynomial of the least number of terms. However, since most error events are bursty in nature this point is perhaps academic, and the simple 'multiply by the number of feedback taps' rule no longer holds. Indeed, a precise analysis requires an exact knowledge of both the error burst characteristic and the state of all elements in the system.

It is possible to remove the adverse state detection mechanism at the descrambler (it must remain at the scrambler). With the example type of adverse state detector shown in Figure 2 (i.e. inversion of the last delay stage only), this leads to only a single binary error at the descrambler for each operation of the adverse state detector at the scrambling end. The removal of the adverse state detector at the descrambler will be examined again later from a consideration of the need for transparent transmission of AIS to the customer.

The bounds on the adverse state detector threshold and its relationship with AIS are considered below. Before any



bounds can be determined for each network layer it is important to understand how AIS manifests itself in networks, which in turn requires an understanding of how the scrambler/descrambler will behave under such network aberrations and what the effects are on the original signal. In turn, this requires a consideration of the parameters associated with AIS events since this is the key to determining what performance measurements are possible.

The way in which AIS manifests itself has already been discussed previously, where it was noted that 'soft failures' can result from error events in server network layers, and that an uncontrolled bit-slip can propagate through the client network layers.

The implications of an uncontrolled bit-slip for an unstructured signal with bit rates at or above the primary network layer of the Plesiochronous Digital Hierarchy (PDH) is that the descrambler will be affected by the extra or deleted bit in a similar fashion to an error event. That is, the descrambler will need to recover synchronisation and the time taken will be proportionate to the number of stages in the descrambler - for the example scrambler/descrambler embodiment given this could be up to 4 clock cycles. In practice, since the uncontrolled bit-slip will invariably be accompanied by surrounding error activity this effect is not that significant. The key point to note however, is that a self-synchronising scrambler/descrambler will 'process out' the uncontrolled bit-slip and recover synchronous operation.

Experimental work has shown that a 2048kbit/s primary rate path signal structured in accordance with ITU Recommendation G. 704 and G. 706 will recover framing and normal operation about 1 to 2ms immediately after the point at which the disturbance which caused loss of framing is removed. It is understood that a longer reframe period (say 3 to 6ms) might apply to the 1544kbit/s primary rate path signal



structured in accordance with ITU Recommendation G. 704. Since the nx64kbit/s network layer is usually synchronous with any serving primary rate network layer, a corresponding period of disturbance will also exist in the nx64kbit/s path signal.

Events of the above type give rise to an effective bit error probability of 0.5 over the period of disturbance, and since they are troublesome to certain customers it is important that they can be measured. Moreover, it is interesting to note that a bit error probability of 0.5 over about 2ms represents a BER (Bit Error Ratio) of about 10^{-5} over 1-second. From a G.821 or G.826 perspective, such an event can be considered to be a SES. Moreover, it is generally accepted by experts in the field of error performance that a loss of framing event at any network layer can be considered to constitute a SES event at that layer and all supported client layers.

When a structured network layer path signal loses frame alignment (or experiences a loss of signal) it generates an AIS downstream as described previously. Since the adverse state detector in the NTU prevents the scrambler from generating a consecutive sequence of more than n (say) binary '1s', then a sequence of greater than n binary '1s' detected at the receiving NTU can be deduced to originate from within the network and hence can be associated with a SES event. The next step, therefore, is to determine a method for setting the value of n at the various network layers at which unstructured services might be offered by network operators.

A scrambler of m stages cannot generate a consecutive sequence of '1s' greater than 2^{m-1} when running through its maximal cycle for a constant input of "0"s or "1"s. If, in the example embodiment of an adverse state mechanism given, the summation device threshold is greater than 2^{m-1} , but less than a larger value n say, then any sequence of greater than



n binary all '1s' - can be associated with an AIS (and hence SES) event from within the serving network infrastructure.

For the example embodiment of the scrambler shown (i.e. 4 stages) this means that the threshold of the adverse state detector mechanism should be set greater than 7. Hence, suppose the adverse state detector threshold was set at 10, then any sequence of greater than 10 consecutive binary '1s' detected at the receiving NTU would be considered to be an AIS (and hence SES) event from within the serving network infrastructure.

The above can, of course, be generalised to other embodiments of scrambler and adverse state detection mechanism as required. However, it is worth noting that bursts of AIS associated with loss of framing at network layers above the primary rate can be quite small, and so a scrambler device should be chosen which does not generate long sequences of '1s' during its maximal cycle, when the input is a continuous sequence of "1"s or "0"s.

In order to make the system flexible so that it can cope with the various path layers at which unstructured services could be offered in different network hierarchies, the value of the threshold at which the adverse state detector operates and the threshold at which an AIS (and hence SES) event is detected at the receiving NTU should both be made independently programmable. The independently programmable range for each threshold type could be different.

From these considerations the definition of a detection criterion for a SES at a receiving NTU which could be used at various network layers could be as follows:

A SES is deemed to have occurred on an unstructured path when n or more consecutive binary 1s are detected in any 1-second period, where n is programmable over the range 10 to 100 (say) in unit steps.



Correspondingly, the threshold at which the adverse state detection mechanism operates at the scrambler 4 and the descrambler 10 could be defined as follows:

The adverse state detection mechanism of the scrambler operates when a consecutive sequence of R binary '1s' is detected going into the network from the transmitting NTU, where R is programmable over the range 2^{m-1} to n (say); where m is the number of stages in the scrambler and n is the independently programmable threshold set for the detection of a SES event at the receiving NTU.

The adverse state detection mechanism of the descrambler operates when a consecutive sequence of R binary '1s' is detected coming into the receiving NTU from the network. R is programmable over the range 2^{m-1} to n (say). Where m is the number of stages in the descrambler and n is the independently programmable threshold set for the detection of a SES event at the receiving NTU.

The adverse state detection threshold set in the transmitting NTU scrambler and in the receiving NTU descrambler should be the same.

A mechanism has been described which resides before the descrambler in the receiving NTU 8 and whose function is to detect SES and AIS. In one example embodiment given previously, if a period of true AIS is input to the descrambler, and the descrambler has an adverse state detection mechanism as shown in Figure 2, then the descrambler will, with period $2^m - 1$, insert a '0' into the output signal transmitted towards the second receiving end of the path 11 to the CPE 12, i.e. there will be a sequence of $2^m - 2$ consecutive '1s' followed by a '0', and this pattern will repeat as long as the AIS persists into the descrambler.

This might or might not be a problem. However, if the customer expects to see the AIS on network failures some



method of allowing the AIS to be output to the customer will be required.

There are several ways this could be achieved, and some examples are as follows:

- i) Removal of the adverse state detection mechanism at the descrambler altogether.
- ii) The detection of AIS at the input to the descrambler disables the adverse detection mechanism for the duration of the AIS. Noting that this could allow one or two operations of the adverse state detector (i.e. insertion of '0s' into the customer signal output with period $2^m - 1$) before it is disabled.
- iii) Detection of AIS at the input to the descrambler causes a 'bypass' function, of both the descrambler and its adverse state detection mechanism, whereby the customer output signal is forced to the all '1s' state for the duration of the AIS.

The obvious merit of the first option is its simplicity and potential cost saving in the equipment and its operation - though in practice this might be trivial. The merit of the other two options is retention of the adverse state detector for correction of customer generated sequences of all '1s'.

A further point which should be noted here is the relative frequency of network-sourced error events compared to the frequency of operation of the adverse state detector at the scrambler. Consider the case where an adverse state detection mechanism was not implemented at the descrambler (noting that it must be retained at the scrambler). It has been noted previously that, for the example embodiment given, this results in a single error at the descrambler output for each operation of the scrambler adverse state detector. However, it should also be noted that errors can occur at the



output from the descrambler due to error events sourced from within the serving network infrastructure 6 in normal operation. If the frequency of errors arising in the serving network is significantly greater than those which would arise due to operation of the adverse detection mechanism at the scrambler, then perhaps this could be a valid argument for removing the adverse state detection mechanism at the descrambler.

The above invention has been described with respect to all network layers which do not have a deterministic frame structure and it will be readily apparent to the person skilled in the art that the implementation of the embodiment is merely given by way of example and that the scrambler/descrambler could have a different number of stages, and the precise nature of the AIS could also be changed.

This Sheet Has Been Replaced



This Sheet Has Been Replaced



CLAIMS

1. A digital communications path having means for monitoring the path, the path having two ends, a first end receiving incoming traffic signals and connected to a scrambler device for scrambling said incoming traffic signals, a transmitter device for transmitting scrambled signals along the path, and a second end at which the transmitted signals are received from the first end, the second end being provided with monitoring means for detecting an Alarm Indication Signal in the transmitted signals and unscrambling means for unscrambling the transmitted signals.
2. A digital communications path according to claim 1 in which the incoming traffic signal is scrambled by a scrambler device, the traffic is transmitted along the path, at the second end an alarm signal detector detects any alarms, and in the event of no alarm signals being detected passes the traffic signals to a descrambler for descrambling, the descrambled signals then passing out of the path.
3. A digital communications path according to claim 2 or 3 in which the traffic signals are not transmitted in a frame structure having traffic and path overhead parts.
4. A digital communications path according to any one of claims 1 to 3 in which the scrambler is a self-synchronous scrambler and scrambles the traffic signals by use of a primitive generator polynomial.
5. - A digital communications path according to Claim 4 in which the scrambler is provided with an adverse state detector, said detector detecting the output from the scrambler and when said output exceeds a predetermined

JFR\GHC\PATENTS\A24680



* C L M S *



threshold number, resetting to zero one of the values of the generator polynomial.

6. A digital communications path according to any preceding claim in which an alarm signal is indicated when the output from the monitoring means exceeds a predetermined threshold number of digital "1"s.

7. A method of monitoring a digital communications path including means for receiving traffic signals to be transmitted along the path at a first end of the path, scrambling the traffic signals by use of a scrambler device, transmitting said scrambled traffic signals along the path, at a second end of the path monitoring the path for an AIS in the scrambled traffic signals, indicating the presence of an AIS in the event of an AIS being detected in the scrambled traffic signals, and unscrambling said received traffic signals by means of an unscrambling device.

8. A method of monitoring a digital communications path according to claim 7 in which the traffic signals are scrambled by a self-synchronous scrambler at a first end of the path and descrambled at the second end by use of descrambling means.

9. A method of monitoring a digital communications path according to either claim 7 or 8 in which a fault on the traffic signal is indicated by use of digital '1s'.

10. A method according to claim 8 in which the scrambler at the first end scrambles the incoming traffic signals, an adverse state detector monitors the scrambled input signals to detect the presence of an adverse state, if an adverse





state is detected the adverse state detector triggers a resetting of the scrambler device.

11. A method according to claim 8, 9 or 10 in which the adverse state comprises a continuous stream of $(n-1)$ digital "1s" where n is the predetermined threshold for triggering an AIS.

12. A method according to claim 10 or 11 in which the monitoring means detect the incoming scrambled traffic signal, and if a continuous stream of n digital '1s' is received which exceeds a predetermined threshold n , indicating an AIS.

13. A method according to claims 11 or 12 in which when the AIS is triggered an automatic connection is made to a remote management centre.

14. A network terminating unit (NTU) for use in a digital communications path having means for receiving an incoming digital traffic signal, scrambler means for scrambling the traffic signal, and transmission means for transmitting the scrambled signal along a path.

15. A network terminating unit (NTU) for use in a digital communications path having means for receiving a scrambled digital signal, means for monitoring said signal for the presence of an AIS, descrambler means for descrambling the scrambled digital signal, and a second output means for outputting an AIS in the event that an AIS is detected by the monitoring means.

16. A network terminating unit according to claim 14 which is also provided with an adverse state detector for detecting a continuous stream of n digital 1s, where n is a



predetermined AIS threshold trigger, and if such a state is detected resetting one or more stages of the scrambler means.

17. A network terminating unit according to any one of claims 14, 15 and 16 in which the scrambler is a self-synchronous scrambler.

18. A network terminating unit according to claim 17 in which the scrambler is a 4 stage scrambler.

This Sheet Has Been Replaced



* C L N S *