Proceedings of the Institute of Acoustics

Some principles for control of entertainment electronics by networks

Allen Mornington-West

Quad Electroacoustics Ltd

0 Introduction

The control entertainment electronics using remote control is hardly new. The obvious examples include tape recorders, light dimmers and video cameras. The development of the voltage controlled amplifier (VCA) lies at the heart of many approaches to the control of complex systems such as audio mixing desks whilst lighting dimmer controls and the video switcher and effects generator have long been placed in places remote from the production gallery where the controls are situated.

The world of entertainment electronics covers a surprisingly wide field. Many of the companies who work within one sector of the overall industry are not fully aware of the activities which are taking place in another. As a consequence there is a large variety of techniques which have been devised in order to carry out remote control. Some of this invention has its motivation in the desire to invent something new in order to avoid borrowing a competitor's technology and some arises through attempts to inveigle competitors in adapting the new scheme. The commercial pressures are clear: it is a company's bounden duty to avoid paying a competitor for intellectual property rights (IPR) as much as it is a target to arrange for competitors to pay for using the in-house developed IPRs. There may be a way out of this matrix of partly compatible techniques but this inevitable change will almost certainly be at the expense of recognising incompatibility with previous approaches.

The major change is an increasing demand for entertainment electronics to be controlled by networks. Protocols which have been adopted, and continue to be popular, are being called into question as they are found wanting against the new demands being made on communication between equipments and their controllers.

It is worth producing a few examples from the wide range of equipment which might be candidates for this control. It ranges from the simple - such as plain audio power amplifiers, light dimmer units, controlling video monitor screens, changing slides in a projector, controlling a video or audio tape machine- through to complex manipulations typical of audio mixing desks, dynamically altering digital audio effects units, three axis pitch yaw and colour controlled lighting units, remote control of studio video camera systems, fully polyphonic music synthesizers and many more. Some of these equipments have also the requirement that their control is tied to a specific event or to some form of time code. Increasingly modern entertainment systems involve complex mixtures of technologies from the full spectrum of activities of companies working in the professional audio, video. lighting control and musical instrument industries (PAVI). At the present time even modest installations will involve the designer and operator in handling three or four mutually incompatible control techniques. Increasingly the equipment is being placed well away from the operator's control surface and yet there is a strong desire, if not an absolute need, to be provided with information which indicates how well the remotely located equipment is functioning.

A further change is that the amount of data and the range and quantity of its sources has become much greater over the past few years. As a consequence nearly all of the remote control techniques in use, whether proprietary or in the public domain, are being found wanting. Replacements for them are being sought and each sector of the industry is presently investigating which technology to adapt or to invent. Almost inevitably there is a tendency to avoid using an approach which looks likely to be adopted by another branch of the PAVI. Partly this is because there do exist reasons for some of the differentiation but this also reflects the very strong motivation provided by the Not-Invented-Here (NIH) attitude.

1 Applications of network control and monitoring

There are many functions which a properly thought through networking system could provide. Simple fixed stage lights can now be a complex unit with a number of electric motors and control systems which may include the electronic ballasts and dimming circuitry. As a further example we could consider the shift of the final gain control in sound systems from the mixing desk to the power amplifier and that, in accordance with good engineering practice, associated systems such as crossover unit and equalisers are also being moved away from the sound balancing position to be closer to the power amplifiers. As a consequence a simple object such as an audio power amplifier is becoming a more complex item of equipment to control and monitor. Current designs on the drawing board show products which use a powerful 16 bit internal microcontroller and the cost of installing the silicon can be brought down to acceptably low levels. Managers now expect high productivity from staff working with modern programming tools and this means that powerful applications can be written and the product brought to market in shorter development times. In addition to the obvious facility of remote control over gain, the network can be used to inform the operator of the local thermal performance, the signal integrity and provide a record of the past history of signals which have been passed through the amplifier. It is possible to monitor the input signal for its spectral content or amplitude and adjust the gain setting automatically if the attached loudspeaker drive unit would run the risk of being driven inappropriately. Further the facility of being able to review the past signal handling history might be of significance to environmental health and safety at work officials.

The use of network control also eases the facility of making the response of equipment dynamically dependent on the environment. Amplifiers powering speakers in a leisure park, for example, may have their gain adjusted according to wind speed and direction in order to reduce noise pollution. A stage lighting follow spot may be arranged to track an artiste's postion on the stage set automatically whilst at the same time arranging that the apparent source of an audio effect is made to follow at the same time.

Recent developments in lighting luminaires and control have produced units in which full control is provided over three axis positioning, rotation, pitch and yaw, iris, focus, gobo selection and rotation speed, colour and, of course, brightness. Though such units are sufficiently expensive to be deployed with care, it is not unusual for a rock concert stage to employ up to 1000 simpler lights over which control of position, brightness and colour may be required. Many of the changes have to be timed to match the music or the event but this can be more than easily achieved by preloading the data relating to the required change and making it active using a global enable command. This is an approach long established in other industrial activities.

A network will be expected to provide an operator or system supervisor with information concerning the well being of the system under control whilst, at the same time, permitting manual adjustment and recovery from failure should this be needed. Control and monitoring over long distances may be required and this too has pushed the viability of many of the control techniques in use in PAVI to its limits and into implementing work-rounds to extend the original communication technique in ways which were not imagined at the outset of its use.

A further input into the equation is that the number of equipments and the complexity of their interconnection is increasing rapidly. This increase is not matched by the rate at which competent operators for this equipment are being trained. Owners of venues and events where such equipment is used can not afford to carry extra staff nor can they bear the impact on business created by gaining a reputation for a less than marketable product.

Owners of venues and riggers setting up systems will also welcome the advantages in simplified installation. In part these come not just from the network control systems where a single cable passes control and monitoring information and which replaces a web of separate control lines form each equipment. In part the simplification comes also from the increasing penetration of multiplexed signal distribution, usually using digitised signals on optical fibre. These replace the heavy and costly bundles of copper based multicore.

The almost universal acceptance of graphic user interfaces (GUI) in personal computers which are used as controllers in many systems indicates that there is growing acceptance of this particular style of man machine interface (MMI). This is reflected in the increasing ease with which GUI based applications can be developed for a particular products and venues and the prevailing view that the software user manual should only be disturbed in times of national emergency. This does not mean that dedicated control surfaces are not being designed as there are very good ergonomic reasons for this individuation. Controls and their monitoring indicators need to be laid out to suit the operator and the equipment being controlled and this is still best done by those whose work is closely coupled to that field.

One of the issues which face all new equipment designers these days is the requirement to demonstrate compliance with electromagnetic compatibility (EMC). Historically many of the companies which are prominent in the industry today are based on the business activity of enthusiastic amateurs. Audio engineers, for example, must now understand that their designs must be characterised over the bandwidth from DC to 1GHz. Networks will need to be characterised as capable of complying with the mandatory emission and immunity limits as a minimum expression of their robustness.

2 Networks and their users

A list of examples where network control and monitoring techniques would be beneficial covers simple installations in a musician's home studio, through the grander scope of a comprehensive home theatre system and on to fully fledged broadcast station control. In the middle there are the examples of leisure parks, virtual reality arcades, night clubs, churches, television theatres, AV presentations and theatre systems. Examples of unglamourous applications include voice evacuation systems and public address systems, railway stations and surveillance.

This array of examples implies that it might be possible to devise a single network concept which could be used throughout. This is unlikely to take place, however, for a number of reasons. The natural reason is based, very often unconsciously, on NIH but there exists a range of considerations based on the kind of objects and the distance over which control is required which may suggest that differntiation is inevitable.

There are obvious examples of controllable properties in each simple item of equipment. A power amplifier would have its gain control adjusted whilst a light dimmer is a clear candidate for remote control. A video camera in a TV studio could be remotely controlled and this is really no different to determining that items of stage equipment could also be similarly controlled. Many of the controlled objects are adjusted only rarely as in the case of video distribution amplifer gain and equalisation. The data rate could be broad and this suggests that a hierarchy of signalling speeds may be justified.

Even if the primary purpose of the network is to control equipment there is stll a need to be able to monitor the performance of the equipment. Increasingly audio power amplifiers are placed closer to arrays of loudspeakers which are flown well out of reach of any casual maintenance access. Further, the failure of any one of the amplifiers or its load may not be acoustically noticeable from the monitoring position but it may affect the quality of the sound received by a section of the paying audience. Similar examples can be drawn from the lighting and video industries. The situation in respect of controlling stage mechanics, either at the coarse level such as opening the stage curtain down to finer levels involved with animatronics, is similar only there may be the added issue of safety.

One obvious example where a feedback information path is usually required is in the control of machines such as video and audio tape recorders and disc storage systems. These are often required to deliver or follow timecode and it is thus important that in such a system that time code signals will be transmitted without being delayed more than a determined amount by other network traffic. This particular requirement sets one of the criteria of any candidate technology and this need for deterministic access is called access latency. By the same token it may be important that other forms of machinery are controlled by this guaranteed delivery of signals. An animatronics model which is forced to move inelegantly due to the hesitant reception of control signals is not acceptable as a paying public has come to expect that the smoothness of movement which is achieved through the skills of the film and television animator are, in fact, capable of being shown in reality.

Monitoring PAVI equipment gives the operator an insight into a number of operational parameters which may be important in determining the optimum preventative maintenance period for an equipment or one of its constituent parts or even in providing some indication of a failure mechanism. Obvious characteristics such as lamp brightness, supply voltage, load current, colour balance in a TV monitor, air temperature are all capable of being used to improve the confidence with which an equipment can be deployed and trusted to behave as required. The record of lamp brightness could be used to warn the operator of impending lamp failure and thus prompt a change during the next maintenance session. More subtle uses can be envisioned, for example, in an audio power amplifier where a local record of the peak voltage delivered by the amplifier over a determined active period can be retained and reviewed later. This information could be used to manage evidence of equipment abuse or to confirm to local environmental health officers that the related sound pressure levels were within agreed bounds. It may be beneficial to permit remotely sited equipments to be adjusted locally by hand to suit a particular requirement and, following satisfactory adjustment the use of the control can be l;ocked out by the network operator thus avoiding the risk of unwanted further afdjustments.

3 The OSI network layer model

There are many remote control techniques in use at the present time and each of them has been devised either to solve a particular set of problems or to avoid previous IPRs or to create the opportunity to charge other users for the right to use a new technique. The term protocol indicates that there is an agreed set of rules used in communicating, usually in both transmit and receive directions, which covers much more than just the determination of the logical high and low voltage or current levels. It will include the ability to handle errors, network, maintenance and transmission using other network systems.

In order to assist with the design of data communication systems the ISO proposed a seven layer model, the Open Systems Interconnection (OSI) model.

1 *Physical*

This is primarily concerned with the transmission of the unstructured bit stream over some physical link. Signalling levels such as voltage, current or radiation are determined along with relevant timing. This is the region where conventional mechanisims such as RS-2332C and RS-485 are defined. Although the OSI model does not define it there is really a layer 0 in which the details of cable and connector usage are set out.

2 *Data link*

This layer determines how the reliable transfer of data across the physical link will be achieved. Thus it detemines the blocking and framing of data along with synchronisation, bit stuffing, channel modulation and flow control. Examples are the packet structure in many protocols in which the error detection power of the frame is usually provided by computing a Cyclic Redundancy Check (CRC). This is calculated by treating the data to be transmitted as one long binary number and dividing this number by a suitable prime number. The remainder is attached to the frame being transmitted. At the receiving end the same calculation is carried out on the received data and the two CRCs are compared. A 16 bit CRC, formed by dividing the data by a 17 bit prime number, can detect all single bit errors, all double bit errors, any odd number of errors and most burst errors.

3 *Network*

This is the layer where connections are established, maintained and terminated. Termination might be characterised, in the simplest case, by response time-out. Problems arising form collisions would also be resolved at this level particularly in conventional computer protocols.

4 *Transport*

This provides the means whereby successive data packets are delivered free of errors and in the correct sequence. In some network protocols some of the error control is handled by layer 2 and much of the maintenance of the connections is handled by layer 3. If the scope of use of a network protocol suitable for PAVI use is opened too wide then this layer would be needed. On the other hand if the relatively local nature of PAVI equipment control is recognised then, provided layers 2 and 3 perform adequately, this layer can become vestigial.

5 *Session*

The session layer provides the services of handling a dialog between two equipments which might arise during the initialisation process for example. It would be implicitly involved in setting up the instance identifier for Convergence Protocol Entities (CPE) and it may also be required to implement parts of a transaction protocol by requesting the full retransmission of a data file should one of the blocks be delivered with errors.

6 *Presentation*
This layer handles the decompression and decryption of data, for example, as it is basically involved in defining the format of the raw data which is exchanged between programs. Within the PAVI the need for this level should be sparse because, in a network whose traffic is composed of many very short messages, there is little need for encryption and neither is there usually sufficient processor power to carry it out. This reflects the need to be able to implement a PAVI network in a most cost effective manner.

7 *Application*
This layer is really the interface to the actual user programs which are being executed. In conventional computer systems the functions of file transfer and terminal access to remote computers would be located here but, within the simpler demands of PAVI network protocols, the simpler functions of local memory backup and input data routing might be all that is required.

Each of these layers group together a sub-set of the functions which are required to communicate with another system in such a way that details within a layer may be changed without requiring consequent changes in other layers. In addition to these layers ISO has also set out standards for network management which may be relevant to controlling PAVI equipment. It is most likely that very few equipments will ever be fitted with excess memory and processing power such that a full protocol, such as the ethernet protocols used in the mainstream of computer networking, will be made available. These higher protocol functions are likely to exist in truncated form only and recognition of this is provided in a subset of the seven layer model in which the upper five levels are collated into a CPE. Within the CPE model an affiliation exists when a remote and local CPEs have knowledge of each other's addresses. This affiliation is used to create an instance identifier or address alias which is unique within the particular network and operating session. The process allows for retention of CPE identifiers for future use as well as their reallocation after replacement of re-initalisation.

4 Forms of network

In devising a network protocol which might suit a range of users in the PAVI the most obvious places to borrow technology from is the computer industry. There are a few terms which describe the basic network architecture used, for example, to link personal computers or industrial process control equipment. In practice any competent practical realisation of one network form should be capable of providing the services which are invoked in another though this may take place with considerable differences in efficiency.

1 *Requester/Server*

Also known as Client/Server this is the principal form seen in systems which are interlinked through the services offered by a file server. In the simplest examples all messages are routed through the server and it is the server which is used to make available frequently used data or programs to users. In the general case each node bears a unique address which is used to communicate with another address and this means that direct node to node communication can be carried out simply. It is not the best model for PAVI equipment partly because of the way in which it is designed to transfer large data and program files. In the practical examples of these networks, such as ethernet based networks, it becomes essential to be able to identify one particular receiving node from a very wide number. It is usually not possible to predetermine all of the states which each node could take partly because there can be no control over the way in which the human user can interact with the node and this imposes a significant extra burden on the software required to implement the network protocol.

2 *Producer/Consumer*

In this form of network there is no *de facto* need for a master node. Any node which has information to transmit does so and all nodes which have a need to use that information are able to receive it. This form of network is very widely used in industrial control networks where, as an example, a temperature sensor located in one part of a process sends out data which is then used by a number of receiving units which adjust process parameters accordingly. In many of these forms of network each node is allocated a fixed node identity and, in the case where replacement is needed, it is sufficient to replace the failed node with a replacement bearing the same identity code. Very often the node identity relates directly to a simple determined set of functions which can be carried out by this node. Initialising such a network often involves manual setting up identities, priorities and programs. Once the use of an identity for a given node has been established for a particular purpose a replacement equipment can be substituted and will perform correctly provided its identity has been correctly provided. Typical in Producer/Consumer models is the understanding that the number of nodes which will be connected to a network branch will be limited and thus, although facilities are usually available for direct addressing, the communication betwen consenting identity codes is adequate.

3 *Teacher/Learner*

Technically this is more of a description of a higher layer of protocol working but it is included here because of its impact on the way in which the detail in the communication between equipments is abstracted. It is also a way of combining some of the properties of the Requester/Server and Producer/Consumer models. The ability to request large data files is restricted and, though this can be done, the overhead is excessive. This is acceptable if the need arises infrequently. The usual form of messages will be relatively short datagrams in which the content and the meaning has been explicitly determined through the use of the Teacher/Learner protocol. The protocol would be fully invoked at first encounter with a previously unknown equipment either through re-initialisation or through addition to the network bus. The number of units which would be attached to each segment of such a network would be limited in order that the overhead invoked by direct addressing was reduced and the problem of addressing more distant units could be handled by extending the CPE identites or aliasses so that messages could be passed over data bridges.

There are also three main topologies for network systems each of which has particular properties. These topologies are quite closely linked to the means whereby equipment nodes access the signalling medium. The signalling medium is usually copper cable in either twisted pair or coaxial form because this still offers a more cost effective solution for local area network schemes than optical fibre or radio.

1 *Bus*

This is perhaps the most common form and it is usually associated with the collision detection method of medium access. Each equipment is connected to the network signalling medium in series (or through a very short **T** junction) and this allows connections between equipments to be 'daisy-chained'. It is usual to arrange that each end of the bus is terminated with an appropriate resistance in order to avoid signal reflections. A signal created at any one place on the bus is able to propagate the full length of the bus. It is normally not possible to bifurcate the bus passively because of signal reflections and hubs, bridges or routers are required to enable both extension of the bus and to increase the numbr of nodes which can be attached to a segment. Removing an equipment node can normally be arranged to avoid disrupting the bus communication. Whilst the bus is being used to conduct communication between two of its nodes it is effectively out of bounds to all other possible nodes. Thus although the signalling speed may be great the rate at which all nodes may be individually accessed is inevitably slowed down.

2 *Star*

In this arrangement each equipment node is driven from a central point or hub. Wiring a rack of equipment will require one network cable for each equipment node reaching back to the hub. The principal disadvantage is thus the cabling but in addition there is the cost of the hub. This is usually an active unit but for practical reasons it may support a limited number of connections and thus a large array of equipment nodes may attract an appreciable charge for providing the necessary hubs. The one benefit which may be obtained, should the hub have sufficient processing power, is that the rate at which individual nodes can be accessed can be very high. However this processing power is associated with a cost and unless all of the hub's ports are used the overhead involved in providing a PAVI equipment node with this form of network control access can be too high.

3 *Ring*

The ring topology is most closely associated with token passing mechanisms for enabling access to the signalling medium. This is not an exclusive property of rings as there are good examples of linear bus topologies which use token passing mechanisms. Physically each equipment node is connected to the next until the connection is returned to the beginning of the ring. The networking signal is arranged to pass through each equipment node where it is processed and, if appropriate, the token is passed on. This places a demand on the processing power at each node which is usually acceptable in computing environments but which will add an excessive cost burden onto PAVI based equipment nodes. One potential problem area exists when equipment nodes are powered down or are removed as the ring may become broken. A variant of the token passing strategy does permit more than one token to be in use at any one time and this can appear to increase the data handling capacity of the overall ring.

5 Key requirements for a PAVI network protocol

Before the key requirements for a network suitable for controlling some, if not most, types of equipment which exist in the PAVI, can be outlined it is necessary to agree on some of the characteristics which these equipments have. At the root of this is the recognition that, although PAVI equipments may based around the computing power of a number of powerful processors, the number of controllable and monitorable objects with which each node may be associated is limited. It is supposed that for practical reasons this number is limited to less than 65535. Thus for equipments which have a need for many more than this number (such equipments will no doubt be invented in the future) an alternative strategy will be needed. One solution is to place a number of logical equipments within the same physical equipment case. Though it is possible to devise methods for handling an arbitrary number of connections and internal registers the network protocol is likely to approach the complexity of those established in the computer industry with the consequence that the solution is likely to become unacceptably expensive. It has to be borne in mind that the resulting equipments are intended to placed on the market at the present time and this determines the extra costs which providing networking control and monitoring access can be acceptable.

There is a need to recognise that the depth of intelligence fitted into PAVI equipment will also cover a wide span. At the simple end there might be the single lighting unit with a self contained dimmer and for this unit the incremental cost of installing network accessibility has to be low. At the more complex scale an audio power amplifier or colour television monitor are examples of equipment where the increment in cost can be more readily accepted since the units may already have been fitted with a supervising micro-controller.

A single stand alone rack of lighting dimmers capable of handling some 255 dimmmer channels may have as a minimum only a single register for the current dimming level. A more full implementation would allow for each channel to be associated with a limited number of cues - for example the past three and the next eight - each of which may be associated with a rate of change setting, a the value for the lamp brightness at the present time, the maximum value of the lamp at full mains power and any limit associated with it. This still comprises fewer than 5000 registers. If the equipment has to be controlled in real time then it is probably necessary to ensure that the value of each dimmer setting can be changed within one video frame (40ms). Assuming a 50% network administration overhead and that the new levels, or the changes in level, can be expressed using a single byte for each channel then a burst data rate of around 80KBaud may be required.

Amongst the more complex lighting units appearing on the market at the current time are those which offer control over position in space in addition to tilt, direction, focus, iris setting, barn door positions, colour in terms of RGB co-ordinates, choice of focal plane slides and internal mirror rotators perhaps some 32 registers but this number will be extended if the lighting unit is sent a list of cues which are executed at the appropriate time or in accordance with a pre-determined program. The data rate required to update this unit directly is quite low at around 10kBaud.7

An audio mixing desk comprising 128 identical channels each with full parametric equalisation (probably 12 controls), full dynamics processing (a further 12 controls), a full compliment of auxiliary sends, fader, associated controls and routing (suppose a further 20 controls) capped with a dynamic metering system for each channel will comprise some 6000 registers. Setting all of these registers (assuming all of the data is byte sized) requires a burst data rate of 1.8MBaud if all of the settings are to be changed at 25Hz. A more practical data rate arises if the metering registers alone are accessed at 25 Hz and the remaining settings are adjusted at 1Hz. The data rate required becomes a more practical 110kBaud.

Modern digital audio and video effects systems may offer an interesting challenge because their controlling system software is capable of reconfiguring the function of the unit in real time. Thus an audio effects processor may change from performing the functions of a 27 band graphic equaliser and carry out those of a four band parametric equaliser on the same audio information. Since it is not possible to insist on *a priori* knowledge of the changes in equipment structure it will be important to consider a protocol which is able to adapt to the changing nature of the equipment to which it is attached. In this example some 162 floating point registers may need to be changed and though this represents a trivial amount of data the change in the function of the unit does require reconfiguring the controller's understanding of the unit. If a Teacher/Learner protocol is used this will cause some dense traffic for around 200ms during the learning process. A more exercising situation arises if the 27 band graphic equaliser is adjusted since all of the 162 registers must be updated at a rate which allows the filtering to change without producing switching noises. Much of the solution in avoiding this problem rests with the skills of the DSP programming and the filter interpolation algorithms used. If the calculation of these coefficients is performed outside the unit by the network controller then the filter values may need to be changed at audio sampling rate of 48kHz. The peak data rate may then exceed 75MBaud.

A simple case is given by the video colour monitor which may have settings for each of, say, eight picture sources, which compris7e brightness, colour balance, gamma, geometry, lift and size. If it is assumed that, in addition to several test patterns which may be selectable, each register is associated with a local control buffer the number of registers is fewer than 200. A similar number of registers might exist in a professional audio power amplifier where access is provided to the channel gain constols, the internal limiter settings, the temperature measurements and history of the amplifier use. In this case access can be very slow.

A video distribution amplifier card may have only two registers, one for gain and one for equalisation setting though it is arguable whether an individual card would benefit form such direct network connection. A better strategy might be to arrange for the card frame to be the network equipment node.

It is natural to concentrate attention on the physical and datalink layers of performance. Indeed many audio and lighting designers can be heard to declaim readily that they have an RS485 (or whatever) network. It is as if the higher layers of protocol are just not important and, perhaps, this is just the reason that the current techniques are coming to the end of their life. This probably establishes the following key criteria for both a technology and its performance.

1 Signalling speed and segment distance
The target is at least 125kBd over 500m
2 Signalling distance
There is usually a compromise to be paid where access to the signalling medium is to be shared and yet messages are intended to be sent at a high rate. In polled systems the operating distance is seldom a problem but such systems require a master unit in order to co-ordinate the polling.
3 Communication between peers as well as master, if present
This permits units to derive some benefits from network accessibility without the presence of a master control unit.
4 Low EMC signature and susceptibility
This aspect is increasingly an essential consideration in the light of the EU EMC Directive which comes into force on 1994 January 01
5 Scalable cost, low cost for simple equipment
A useful target is for the incremental cost of adding network accessibility to be kept below 5£. This does assume that the unit justified some form of microcontroller initially.
6 Scalable discovery method, this is the Teacher/Learner protocol

In a low cost equipment it has to be possible to implement this algorithm simply.

7 High reliability
This expresses the need to be able to rely on secure communication and it has an impact on the options which can be made available for the physical layer. Error correction should be built into the network protocol from a very low level upwards.

8 Simple installation costs
The cabling should be readily obtainable and simple to install. It helps if there are LSF forms of an acceptable cable so that its use in buildings can be approved. The connector forms chosen should be reliable and mechanically tough.

9 Isolatable
This is the requirement to achieve galvanic isolation. This is important in much of PAVI equipment as separate electrical supply phases are common in full systems and local earths may be at differing potentials.

10 Extendable with bridges and routers
This is probably one the acid tests of a full network protocol. This feature has been developed for some network protocols whilst for others there are hooks available for its later introduction should the need arise.

11 High number of nodes which can exist on any one segment.
A segment is a section of network which is attached to the same signalling medium. Most practical systems have a limitation because of the loading caused by the number of receivers attached to the bus.

12 Bounded latency
This feature is needed because certain applications (those involving time code or stage cues are examples) where a control message (or a warning message) must be guaranteed a delivery within a specified time.

13 Separation between control and program signals
It is desirable for the program material and the control and monitoring signals to be able to arrive separately at the destination. In a few circumstances it may be convenient to use the same carrier for both classes of signal but this transport method is probably limited to one direction in many cases and is probably unworkable as soon as the number of independent program sources exceeds one stereo audio channel or one colour video channel. There is no ready parallel in the case of lighting or musical instrument control.

A practical networking technology would ideally be in existence already and all that would be necessary is to retrieve it from wherever it is currently in obscure use. The PAVI is an industry is characterised by many small companies (and a few very large ones) who are not used to commissioning their own silicon chip solution to a problem and this is especially so where there is no industry agreed standard. Two examples show that it is possible and one is the production of SMPTE time code generator and decoder chips and another is the AES/EBU digital audio interface. However the idea that there might be one signalling standard for widespread use throughout PAVI is probably impractical partly because there is a large installed base of equipment using older signalling methods. This installed base will have to adapt to the new technology because newer equipments are able to provide more flexibility to the user than the old and amongst these benefits is the power of network control.

Some key requirements of a candidate technology are:

1. *Open technology*
   Companies who are attempting to market their signalling technology try to avoid allowing a technology becoming open as it usually means that there is very little money to be made from the rights to use the technology. It does, however, shift the emphasis away from earning money by selling the technology to making money by selling the equipment which uses it and this, in turn, focusses the design on to doing that job better than the competition. Most of the successful and widely adopted protocols are open and include examples such as ethernet, AES/EBU and MIDI.
2. *Licence free including runtime licence*
   This also represents a problem to companies who wish to dominate the market by charging for runtime licences for software developed for a particular technology. However, since the PAVI is comprised of a great number of small companies it seems important that the costs of producing small quantities of equipment, usually to satisfy special requirements, should be capable of being well controlled. A licence free technology does that.
3. *Multivendor silicon should be available*
   The practical requirement for an equipment node to service network interrupts could absorb most of a processor's time unless dedicated interface silicon is made available. Multivendor availability implies that there is also more than one level of providing compatibility with the network technology and this may help to the provide scalable costs referred to above.
4. *Development tools and support*
   The existence of practical and affordable development tools is important if a reliable implementaiton is to be made and the time between the start of design and the marketing launch is to be kept short. Some of the network technologies have been implemented around popular microcontroller cores and this makes available cost effective real time operting systems which may be required by some applications.
5. *Standards co-ordination*
   Where a disparate band of users are gathered together it is important that there is some mechanism or organisation by which developments and enhancements to a networking technology can be agreed.

5 Existing technologies

There is a great number of communication and networking technologies in current use of which only a few are presently in significant use in the PAVI. This is not an exhaustive list and neither is it a comprehensive review but it covers those most likely to be considered for use by the PAVI. Should none of them be considered suitable then there may be no alternative but to invent someting new but engineering has been described as the art of technological compromise.

1 *RS-232*

This is same signalling approach which is found in nearly all computer serial ports but it has been widely adopted as the specification of the signalling technique in many so-called protocols. It was originally intended for data transmission over distances up to 15m at speeds up to 9600 Baud but it is stretched well beyond this speed and distance by many users. It is not a balanced system and it is not even specified to be used with a screened cable so that, these days, its robustness in EMC terms is open to question. Indeed the long length of unscreened cable is probably a very good radiating aerial in many systems.

2 *RS-422 and RS-485*

Both of these standards refer to balanced voltage drive systems. RS-485 is intended to be used in segments where up to 32 transmit and receive buffers may be connected to the same bus segment. The specification recognises the importance of using matched cable and driving impedances and it can, in principle, be used at rates of 10kBaud over 2km.

3 *Lonworks*

This is a proprietary offering from the Echellon Corporation which is based on the token bus concept. It has been used in a wide range of applications many of which are in the field of building monitoring where it appears to have a better market share than the european BatiBUS protocol. Silicon is available from two vendors and there are licencing costs to consider.

4 *RC5*

The world of domestic products has long had a simple standard for one way remote control based on the use of infra red transmission. The standard specifies exactly what each transmitted code shall do for each type of product and in that regard it is quite fixed. Since there is no reverse channel it can not really be considered as a network standard. Silicon encoders and decoders are available.

5 *ES Bus*

The European Broadcasting Union and the SMPTE conspired to produce a signalling system which would allow the control of video and audio tape recording machines. The standard is based on RS485 drivers signalling at 38.4kBaud and has a considerable amount of its specification relating to handling SMPTE time code and to consequent events. For the most part the framing of each character is accomplished using standard UARTS.

6 *PA-422*

This is an early attempt by the Audio Engineering Society (AES) to determine a standard for audio electronic equipment control. It is based on the RS-422 signalling system but it has not been widely adopted partly, perhaps, because it was not possible to implement peer to peer communcation and also perhaps because the definition of each of its messages was too tightly defined to allow for easy expansion. There was no need to specify special silicon since the driver chips and the UARTS remain standard parts used throughout the computer industry.

7 *D2B aka IEC1030*

This is also a standard which was intended for domestic product use. The standard accommodates two directional signalling but it has not seen widespread use in consumer products partly because it is quite expensive to design and manufacture a handset for the cost sensitive consumer market.

8 *DMX-512*

This is the protocol which is used throughout much of the stage lighting world. it is based on using the standard UART at 250kBaud with RS-485 line drivers and, without a repeater it provides a reliable signal up to 250m. There is no error checking on any of the data. The original standard specified a twin pair cable in which the second pair could be used, at a later date, to enable a reverse channel. This is rarely achieved. The standard was originally intended for the remote control of racks of remotely located dimmers but it has been adapted, through redefinition of the header or channel data start byte, to provide a means of controlling other equipment in the realm of stage lighting, effects and stage mechanics. It is not an elegant approach but it does have a wide installed base despite the fact that the signalling system can only be galvanically isolated by using low threshold opto isolators connected directly to the signalling bus. It is actually ripe for thorough revision but there are problems with trying to reconcile backwards compatibility with the need to produce a more flexible system. Although the signalling system has the capacity to update the dimmer information for each of up to 512 channels at approximately 40Hz it is rarely achieved. This is because of the surprisingly large amount of data processing which modern lighting control desks are required to perform. A more elegant approach would produce the command bitstream on a number of network channels some of which would then be used to control electromechanical systems in the lighting units themselves.

9 *Ethernet*

Ethernet in this context usually implies the linear ethernet bus which uses the Collision Sense Medium Access Collision Detect (CSMA/CD) bus access and arbitration method. Most new installations employ the twin twisted pair form of ethernet in which one pair carries the received data from the bus and the other the transmitted data to the bus. In this guise it is limited to a spur length of around 100m. Although the signalling rate is 10MHz the data rate is half of this since because of the Manchester coding. The uncontrolled bus access time caused by CSMA/CD rules out its use for guaranteed machine and event control. Further, the minimum data packet is 46bytes and this is an addition to the high overhead which has to be supported when transmitting a number of short messages. There are many sources of silicon but most have been optimised for easy interfacing to personal computers and they are thus not a natural fit into PAVI equipment. Further there is a significant overhead involved in handling the various messges which the interface chips must provide and it is doubtful if a cost effective interface can be achieved.

10 *Crown IQ*

This a proprietary system whch uses RS 485 drivers. It is in limited use.

11 *Nexys*

This is another proprietary system which uses RS-485 and is in limited use.

12 *SCSI*

This is the small computer systems interface which has been ued to control lighting systems but it can not be called a networking system as it has severe limitations in the distance over which the control and data signals can be sent.

13 *TCP/IP*

This is the transmisson control protocol + Internet protocol used to establish and maintain communication between mainframe computers. It is open but it also involves a significant overhead to implement and thus it does not appear to be the prime choice for control of PAVI equipment. It is rumoured that SMPTE have considered that it might be the correct protocol for the remote control and monitoring of broadcast stations but the impact of its overhead may have been overlooked.

14 *GPIB*

GPIB is a byte wide parallel communication protocol which has a considerable installed base in instrumentation. Its use for control of PAVI equipment is unknown and it can hardly be recommended as it has a limited signalling distance and a high cost of implementation.

15 *Bec*

This is a proprietary system with limited penetration in the general PAVI.

16 *ARTI*

This appears to be similar in some respects to the ES Bus technology but it is proprietary and its penetration in the PAVI is not great.

17 *ATM*

The future holds that Asynchronous Transfer Mode will become the message carrying system of the future. ATM is based on a fixed packet size which is routed through a switcher at 155MBaud. Thus it becomes possible to consider delivering program material, either audio or video (usually compressed) using this approach. Throughout this discussion on network protocols it has been assumed that the program material will be carried on some independent medium and that the control and signalling technology can be separated from it.

18 *MindNet*

Intellix Corporation's proprietary approach which has been proposed for a number of applications but it is not in wide use in the PAVI.

19 *Media Link*

This is a proprietary bus technology owned the the Lone Wolf Coporation. Licence, product validation fees and royalites are involved. The bus uses a form of token passing and is capable of the 125kBaud signalling requirement. Lone Wolf have also developed some of the higher layers beyond the OSI model layer 7.

20 *Profibus*

This is an industrial control bus with a number of adherents in industry. There is a number of industrial control protocols some of which have had a limited success. Mostly they teach us that it is the open and non-proprietary control technologies which see widespread application.

21 *Arcnet*

This is a potentially useful bus based on token passing using a signalling speed up to 5MHz. It is typically characterised as being able to drive up to 10 device nodes in a 130m segment when using RS485 drivers to drive typical cable samples. At lower speeds longer distances can be readily achieved. It does have silicon and it does have a practical packet size and,. because of its lower packet overhead and deterministic signalling it can convey more single byte messages than the standard ethernet. But it does use 11 signalling symbols to convey each byte and this, in conjunction with the token passing overhead, slows the system markedly if it is used for very short messages. In addition to this a succession of short messages from one device to another can be delayed significantly if there are other nodes on the bus as the token has to be passed to each before it can be used by the original communicating pair.

22 *SMX*

This was a proposal forwarded by Strand Lighting UK for controlling theatrical and broadcast lighting systems. It was not adopted by the industry partly because of its proprietary origin but also because many manufacturers perceived the proposal to be too complicated for them to implement. This is not an acceptable comment since the protocol document is clearly set out and it does attempt to provide, for the first time in lighting system control, a true protocol.

23 *MIDI*

The Musical Instrument Digital Interface is based on the use, or should it be the abuse?, of the standard UART clocked to provide a data rate of 31.25kBaud. From the outset the standard has always required optical isolation in order to avoid hum and earth loops when equipments were connected. Because of the implementation there is a limitation in the distance over which the signalling can take place and there is also a limitation in the number of individually addressable equipments which can be accessed. Since it was conceived as a comand only protocol there was never any provision for a reverse channel. Recently it has been suggested that the MIDI signal could be returned back to the orignating equipment in order to provide a monitoring path. This proposal serves to show how this very popular interface has reached its logical bound. Because some personal computers have always been fitted wwith MIDI interfaces there is a generation of homespun engineers who have used the interface to exercise control over non musical items. Thus there is a range of MIDI like systems such as MIDI show control (MSC) (which can be characterised to handle a full range of theatrical mechanics including pyrotechnics), MIDI machine control (MMC)(which is used to control audio disk recorders and video machines in editing suites) and MIDI PLC which is an interface between MIDI and the ranges of Programmable Logic Controllers which can often be fiund in use in industrial process control. The use of MSC to control lighting would indicate that the raw data rate of MX512 is not really needed in many cases.

24 AES/EBU

This is not a control network technology but a means of delivering digital audio program signals. It does, however carry two user definable bitstreams running at the same rate as the sampling frequency, typically 48kHz, and these could be used for remote machine control. A separate path would have to be provided for \ reverse channel.

25 CAN

Controller Area Networks is a member of the control technologies which were intended for one market and which has seen extensive use in a wide range of other markets. Originally intended for the automobile market (where it is now in extensive use) it has been adopted for industrial process control and, to this end, there is a very comprehensive documented protocol which is readily available. CAN may not be ideal but it does offer 125kBaud signalling over 500m and, despite its small packet size, it is possible to design bridges and routers. These depend on using the CPE approach to local address aliasses. Higher signalling speeds, up to 1MBd over 80m are possible. It satisfies all of the network criteria listed above including two methods of achieving galvanic isolation, some six vendors of silicon parts and an ability (using the proper CAN driver) to sustain up to 110 nodes on a single segment. Although it is originally intended to perform as a producer/consumer network model it can be adapted to a compromise whereby each packet is addressed The address of the sender is either not relevant or it can be inferred from the context. This reflects the basic system requirement for controlling and monitoring PAVI equipment and in this regard it highlights a difference between control of PAVI equipment over a network and the use of a network within the more conventional computer industry.

26 ZIPI

This is a protocol which has been recently proposed as a replacement for MIDI. It too is based on the ubiquitous UART only in this instance it intended to clock the system to provide a 250kBaud burst data rate.

Comparison between these systems is invidious. In the absence of other external factors the practical choice may be limited to those systems which can signal at an acceptable speed over the distance, are open and licence free and have a true bidirectional (or reverse channel) protocol built on them. Those falling through this filter include Ethernet, Arcnet and CAN.

One of the characteristics which separates the requirements of PAVI equipment networking from conventional computer approaches is the messaging speed. One indication of this is to compute the number of 1 byte messages which each system can convey in a second. This comparison can be made easier if it is assumed that the competing systems are adjusted so that they use the same basic signalling rate. This will be fixed at 125kBaud. This performance element can be compared to the time taken by each network technology to move a 125kByte file (1Mbit).

All signalling adjusted to 125kBaud
ethernet limited to 60% bus saturation
arcnet assumed to have only two participants
DMX512 frame configured to controlonly one dimmer channel

network type	*maximum clock speed MBaud*	*number of 1 byte messages/sec at 125kBd*	*time to move 1Mbit at speed*	
			full	*125kBd*
ethernet	10	66	0.2	1.6
arcnet	5	811	0.28	1.13
CAN	1	2717	1.6	12.7
DMX512	0.25	2673	44.3	8.9
MIDI	0.031	2840	103	25.7

These comparisons would appear to be contrary to expectations but this is only because the overhead in the signalling process has been taken into account. Where the overhead in each packet is high it is usual for the packet to capable of carrying a large amount of data. Ethernet loses apparent speed because the Manchester coding effectively halves the data rate. It is clearly not sensible to use CAN to transmit pages of text. The comparison is not fully fair because neither the DMX nor the MIDI protocols are truly bidirectional and they do not take full account of the added overhead involved in handling messages which exceed the packet length. The situation for PAVI equipment is very much oriented towards the sending, and reception, of short messages and thus optimising the network signalling technology to achieve performance in this area would appear to be the right strategy. If a network protocol for this industry were being constructed from the ground up it might be sensible to consider devising a protocol in which the overhead is a function of the message types to be sent.

In the case of CAN there is a penalty to pay for this and this arises out of the terseness of the CAN messaging. Implementing devices such as bridges will require longer messages than a single packet if only to provide a full address to the hub or bridge and this will take time. A significant further time will be required for the discovery and Teacher/Learner protocol to function. Fortunately the Teacher/Learner phase is only required to be repeated once in a session and in the event of a unit being removed and returned a competent Learner equipment will have no difficulty in establishing quickly that the new unit on the scene has been encountered before and can be assigned a valid identity with very little delay.

6 The higher layers

In this text the details of the protocols involved in implementing the five higher layers are left to one side. That is not because they are not important but because there exists a very good body of work in which most of the required functions is described. This document, which is freely available, is the definition of DeviceNet and this is a specific development of the use of CAN for industrial control purposes. Since the document is open it is proposed to use much of what has already been shown to work and to adapt it for use in a PAVI network protocol.

At this stage an army of computer programmers will usually appear and claim that everything they write will be object oriented. Objects will be placed into 'containers' and containers can be nested. Objects may also inherit properties from other objects or may have new properties appended. All this is true but it is detail which has no need to appear in a protocol. What does need to appear is the outline of a Teacher/Learner protocol (TLP). The AES 24-ID is a draft document which sets out to describe a protocol for controlling audio equipment (only) but a considerable part of its thrust is focused on the application of object oriented programming techniques and very little effort has been spent on real protocol.

The TLP approach is a proposal which is intended to bypass one of the problems in exercising network control over equipment which is essentially unknown to a controlling equipment. It would be easy to determine today that, for example, all audio power amplifiers will comprise two channels, each with a gain control, one internal temperature sensor and a warning signal for channel failure. This definition of an amplifier changes overnight if a designer decides that the amplifier should incorporate an overload limiting circuit whose threshold can be altered. In other words it is never going to be possible to fix the control appearance of PAVI equipment. Of course, to the designer of a new item of equipment such equipment may be well defined and to this end the manufacturer will probably devise a custom control surface in order to show off its best features.

7 The Teacher Learner protocol and FLORIES

If the detail of the control is indeterminate or is retained as some kind of commercial secret then it is unlikely that the equipment will be easily attached to some other control network. Some concept of a uniform interface is thus needed to make the best of the chosen signalling protocol discussed above. The starting point is to consider PAVI equipment as being either a Teacher or a Learner. Teachers will, in general, become equipments which are intended to be controlled and may be infrequent senders of information. Learners will often be controlling devices. Of course it is possible for an equipment to be both a Learner and a Teacher just as it is possible for an equipment to have neither Teaching nor Learning skills. An equipment with none or vestigial skills would be very cheap to devise and it would not break the protocol.

To help with the process we will build up a model of the equipment which is to be monitored or controlled by asking appropriate questions of it. Thus, at this stage there is no need to know that the equipment is a lighting unit, a remote TV camera or a power amplifier. Each controllable and monitorable object which it has will be identified with a register. Every control, level adjustment, switch, register is identified by its register number. A further question will reveal the names given to these registers, the size and type of the data which is normally placed in them. The next level of query will reveal how this data is to be manipulated and finally we can determine how the information is to be displayed. This dialogue is intended to be verbose and the data which is communicated should be in english and not a string of indecipherable hexadecimal numbers. As a consequence this discovery or learning process will take time. For example a modern audio power amplifier with some 300 registers may use some 200ms to teach a good learner what it has to offer if the signalling protocol is based on CAN.

The information which is communicated includes details of the manufacturer, model, serial number and class of equipment. This data is used to create a CRC which is used as a shorthand to identify the equipment when it is encountered in a later session. If the learner follows a simple model and loses the data it has acquired when it is powered down then the full learning sequence will need to be progressed at the start of each session. A further few registers are retained for data which should be present in every equipment such as the software version number. In addition a manufacture may wish to log the date of manufacture or the name of the purchaser in a similar manner.

The acronym for the format of these descriptive statements is FLORIES (Formal Language for Object Recognition in Integrated Enterntainment Systems). These statements can be written out in a manner not unlike the WIN.INI file statements in the Windows operating system. FLORIES allows for a few default values for both the type of register (there are only two types potentiometer and switch) and the type of operation which is carried out on it (default types include potentiometer, level, position &c). Data which is placed in it are one of the four types, short, long, ascii and float. Limitations are built into the implementation of any language and, at the present time FLORIES is able to access a maximum of 65000 registers in any one equipment. Learning about the details of this number (probably due to learning about a 128*128 crosspoint matrix switcher) could take a long time but it is perfectly feasible to provide an intelligent controller, which by now is almost certainly based on a personal computer, with the appropriate disc file with the relevant data.

There are some benefits. A FLORIES compatible equipment is able to advise its controller of changes in its internal configuration and this feature would be used by a digital signal processing based unit whose function is changed. A network comprised of FLORIES compatible equipment can include lighting, audio and video equipment all of which could be controlled form a general control panel.

8 Some concluding remarks

What the FLORIES approach offers is the benefit of not having to define all controllable equipment or its combinations today. Flories requires that intelligence is distributed around the equipments and forces the information on the network bus to have a meaning which is only applicable between the two equipments intercommunicating at the time. It capitalises on the ability of a suitable network messaging protocol to be optimised for control at the expense of passing literature between equipment nodes. It is quite well suited to CAN but this is not a unique proposition as there are other signalling technologies. FLORIES and the Teacher Learner protocol allows new equipments to be attached to a control bus and to be placed under control of a suitable Learner. The Learner and Teacher algorithms intrinsically support the requirement that peer to peer operation should be readily achieved. Equipments which have features whch can only be properly invoked by powerful controllers can be readily arranged not to teach simple learners and thus they protect the equipment from potential abuse.