WORKING NOTE AC80

Alan Creak
26 August 1991

CONTRIBUTIONS TO THE DESIGN OF A DOMESTIC COMMUNICATIONS NETWORK FOR ENVIRONMENT CONTROL.

This note presents my attempt at analysing the conditions which constrain Mike Diack's network, which he is constructing as his M.Sc. thesis work. I record the fact that Mike has already designed the essentials of the network; my analysis is intended as no more than a systematic check to see whether he's missed anything important.

REQUIREMENTS.

Aim : To develop a communications network which can be used for control and monitoring purposes in a domestic environment. It should provide access from a control terminal to devices which can be controlled by electronic means, and to devices which can provide useful information.

Domestic networks designed for control are available, and details of a "do-it-yourself" model have been published[1]; a complementary device to permit control through the telephone is also described[2]. While such systems are useful, they do not provide for any return communication to the control terminal. It is therefore not possible to receive feedback from the controlled equipment to the control terminal, nor to use the control terminal as an information or warning device in case of detected events - cooking or washing completed, intruders, fire, or whatever.

Why : Now that it is possible to control a wide variety of devices from afar, quite a lot of people would like to do so - whether because they are physically unable to reach the devices on account of disability or distance, or because they are too idle to make the effort. So far as my own interests are concerned, those three groups are listed in order of decreasing importance. People with disabilities could benefit particularly from the provision for communication from the network to the control terminal. Two examples :

- People who have difficulty in moving around the house quickly, or at all, can use it to check on the condition of various devices without the labour of going to check. Is the door locked ? Did I leave the light on ? Is there anything in the mailbox ?

- Many devices (doorbells, telephones, timers) are designed to draw attention to themselves by sound, which is inconvenient for deaf people; statagems for using flashing lights or other means are possible, but limited in the number of conditions they can easily communicate, and need special wiring. A single device which could draw attention to itself by vibration or physical contact and then relay an appropriate message in words would have obvious advantages.

The proposed system is not a panacea. It is a simple system, with a comparatively narrow data channel, and is thereby precluded from any task requiring high data transmission rates. For example, most people rely heavily on vision in domestic - as, of course, they do in other - activities, either to report measurements (clocks, thermometers) or to give feedback on the results of actions (is the window shut ? has the cat eaten its meal ?), and this sort of communication is difficult for blind people. Despite a formal dualism between converting sound into vision for people who can't hear and converting vision into sound for people who can't see, the data rates involved are very different. The requirements may well be beyond the capacity of the sort of network proposed, unless either it is extended far beyond our present intentions, or the message to be conveyed can be isolated from the visual carrier. It would not be particularly difficult to replace the visual output on the control terminal by a speech synthesiser, and that could be a useful alternative design for the other operations, but the complexity of the information carried by visual cues means that this simple network is unlikely to be capable of handling it in any directly useful way.

What : We don't want to impose any unavoidable restriction on the sorts of device which can be handled by the control system, so we'll try to cover all possibilities. That means that there are about three classes of device : stupid, intelligent, and C-.

- **Stupid devices** can only be switched on and off; so a standard on-off switch will be prominent among the units controlled.

- **Intelligent devices** will probably contain their own computer control units, and be designed for external control. They will have well defined control signals and responses, and probably include their own provision for fault detection and reporting. No commercial examples come rushing to mind, but there are reports of experimental systems incorporating mobile robots[3], and vacuum cleaners have been proposed[4].

- **C- devices** are electrical machines controllable above the on-off level, but they require manual operation and have no provision for control by wire : radio and television, dishwashers, washing machines, etc. These are perhaps the hardest; indeed, it's far from obvious that anything other than a button-pushing robot[5] will do as a general solution. Of course, special solutions are possible : the most common are the remote control units for television sets. So far as possible, though, we would like to provide a solution which can be used with as wide a range of devices as possible without requiring additional development work.

SPECIFICATION.

ESSENTIAL FEATURES :

The system must be as cheap as possible. I won't elaborate this condition on its own in the discussion, but it must be understood as underlying all decisions.

The system must be usable throughout the house but be guaranteed not to interfere with a similar system next door.

The system must be readily extensible to cope with any number of devices.

It should be possible to attach any device (within reason) to the system using only standard system components.

It must be possible to send control signals to the devices, and to receive signals from the devices.

The control terminal must be mobile (so that someone can carry it about).

It must be possible for any device to initiate a network transaction, to permit unsolicited warning messages to the control terminal. (That doesn't preclude a polling organisation, but it means that if polling is used it has to go on all the time.)

Devices may be static or mobile.

DESIRABLE FEATURES :

Possible extension to multiple control terminals. (Perhaps we'd want to cope with input from both terminal and telephone anyway, so this may not be a big problem.)

Possible extension to the garden ? (To control lawnmowers, etc.)

The system should be robust, detecting and reporting failures wherever possible.

ANALYSIS : LOOKING FOR A POSSIBLE SOLUTION.

Constraint : avoid laying new wires (expense, inconvenience), clutter from wires (disability).

For the first constraint : USE EXISTING WIRING or USE A WIRELESS SYSTEM.

WIRELESS SYSTEM :

Electromagnetic, sound : radio the only sensible answer.

For : Passes through walls, low power possible (for portable units), not offensive, not dangerous.

Against : difficult to limit precisely, so could get interference - though by using different frequencies, etc., can get some protection. (Compare door openers - use numeric codes.)

Legal position ?

Satisfies both constraints.

OVERALL : does everything except guarantee non-interference with other devices.

EXISTING WIRING :

Only mains generally available.

For : Passes through walls, not offensive, not dangerous. Connects to most things that would need controlling.

Against : difficult to limit precisely, but probably less prone to leakage than radio Needs special interfaces to get information to and from mains wires.

Legal position ?

Satisfies only one constraint : doesn't provide a communication path to portable units.

OVERALL : does everything except communicate with units not connected to the mains power.

Everything is covered somewhere, but not all at once. Perhaps a compromise ? But only one compromise is possible : if we start with a wireless system, there's no obvious way to add existing wiring to ensure non-interference. (One could add new wires to construct a large Faraday cage around the house ...)

But we *can* consider starting with a mains wiring system, and adding wireless components to cater for portable units.

WIRELESS SYSTEM FOR PORTABLE UNITS :

Electromagnetic, sound : radio no use, because that's the source of the interference problems. But now we don't need to go through walls, so we could use other sorts of radiation within rooms - ultrasonic or infrared.

For : Can be isolated within the room (of the two, ultrasonic is more likely to leak out of the room (it may also terrify dogs, etc.); consider only infrared henceforth), low power possible (for portable units), not offensive, not dangerous.

Against : requires a receiver-transmitter in each room to link the mobile device to the control system.

Legal position not obviously a difficulty.

That should work. Notice that it may be necessary to devise something like a cellphone system if a mobile unit is to work while going from one room to another. While it's straightforward to broadcast all messages from all infrared transmitters, some care may be needed if the same message is received by two or more receivers.

CONCLUSION : USE EXISTING MAINS WIRING, AUGMENTED BY INFRARED LINKS TO MOBILE UNITS.

REVIEW : DOES THE SUGGESTED SYSTEM FIT THE SPECIFICATION ?

THE ESSENTIAL FEATURES.

The system will be able to **communicate throughout the house** provided that there are no "black spots" where the mobile ends of the infrared links can't see the fixed ends. Clearly there must be at least one fixed end in every room, so that we can shut the doors - perhaps as an adaptor between the existing light socket and the light ? Dark corners may not matter too much - Mike's experiments show that scattered light is remarkably reliable indoors as a communications medium - but only experiment will tell.

I don't know about **interference with next door**.

There is nothing in the system so far defined to hinder the development of an **extensible system** - nor to encourage one, for that matter. That's a "higher level" property which must be addressed once the basic communications media are established.

Similar remarks apply for the requirement that the system be **usable with any device**. I say more about this later.

To **exchange information** as required, two issues must be addressed : it must be possible to direct a message to a specific destination; and we must be able to communicate the required information for useful control at an adequate rate.

To guarantee that a message reaches a specific destination, and affects no other destination, we can either use a separate channel (wire or broadcast frequency, for example) for each communication path we want, or we can use a single broadcast system and label each message we send with its destination. The second is obviously preferable on the grounds of simplicity. Apart from that, we've already restricted ourselves to a very constrained hardware channel (the mains wiring), which is not so far as I know particularly well adapted for carrying broadband communications. We continue, therefore, remembering that each message must contain an indication of its destination.

To work out the implications of the need for adequately fast communication, we need a lot of information which we don't have. We may need to know how many units are likely to be used, how often they will need to communicate, the size of the messages, and how reliable the transmission will be. Some guesswork :

- We want to cope with "any number of devices"; but there is a practical limit on how many one can sensibly use. We may wish to use devices for :

  comfort - lights, heating;
  cooking - stove, food processor;
  cleaning - vacuum cleaners, washing machines;
  general housework - robots;
  entertainment - radio, television;
  communications - telephone;
  security - locks, alarms;
  miscellaneous - a safety factor.

  Suppose we allow ten devices in each of these categories (well, can *you* think of ten devices in each category ?); that gives us 80 devices.

- Looking through the list, we see that very few of the devices need anything resembling regular attention in normal use. Rather, they are used now and then - say, a few times each hour at the most. Exceptions are devices which require an approximation to continuous control; vacuum cleaner, robot, and food processor could come into this category. The information requirements of these devices will therefore determine the information rate required. Perhaps a few messages per second will be adequate - say, five messages per second. Provided that each message is in some sense complete in itself (start, stop, faster, slower, left, right, obstacle found, too hot, etc.), and each is a control or information signals between the device and the person controlling it, that message rate is in any case pushing the limits of the human capacity for sending and receiving messages and reacting to them, so should be safe enough. I shall take this rate as a working hypothesis.

  (An alternative way to estimate the maximum load would be to derive it from the human limit. This approach has the advantage that it applies to any number of machines, though it may be called into question by a clever interface device with a large buffer capacity. That doesn't matter much, though, as it's quite unlikely that anyone would wish to work on two of those continuous tasks at once.)

- The messages are likely to be small. Suppose each contains a device address (which, as we expect no more than 80 devices, fits comfortably in a byte), two bytes of information, and a check byte : that's four bytes. Two bytes of information should be sufficient, as the messages are only required to convey simple instructions to the devices (see the list above, and perhaps add a modifier : faster by so much; turn through such an angle), and to return simple items of information (originating device, and result code). (This structure is quite like that suggested for a standard wheelchair interface[6]. The proposed standard message was five bytes long, but included two data bytes - the two orthogonal positions of an analogue joystick - in each frame.)

- How reliable will the transmission be ? If it isn't good enough to carry a good proportion of the messages correctly then it's not much use at all - so I'll assume as a worst case estimate that each message must be transmitted twice, with a "please repeat" response from the addressed device, thus in effect tripling the information rate.

Putting all that together, we get a required data rate of :

(5 messages per second) ×
(32 bits per message) ×
(a factor of 3 for bad transmissions),

giving 480 bits per second. That doesn't look too demanding. Even allowing that it estimates only the useful information rate, taking no account of housekeeping details (start and stop bits, etc.) and possible retransmissions because of network contention, if we can do the job at all it should be possible to reach the required rate. I think.

Finally, I shall address the real demand on the network. Leaving the real discussion to the end like this is not a clever literary or debating device of some sort : I've only just remembered it. I'll leave the preceding analysis there anyway because it gives some idea of what might be termed the background load. The real demand is imposed by the need to get strings of text to the control terminal at a reasonable speed, which is about our only requirement for an extended transmission at a perceptible rate. I've only just got round to it, because it's very much one of the higher level factors; until you begin to think about just what you want to present at the terminal and how you want to get it there, this question isn't forced on your attention. Indeed, I'm only including it at all because of its relevance to the lower level argument I've just given; in this note I never get as far as analysing the higher level details.

How fast do we need to transmit text ? I don't know. People have certainly measured reading speeds, and I recall some work somewhere which suggested that it was better to display text perceptibly letter by letter than to flash a whole message on at once; but I don't remember any details. For the moment, I'll guess that if you can display a line of 80 characters in a second, that's adequate. Well, so much for the 480 bits per second. (Six-bit characters used to exist and worked all right; but we don't want to mess about packing and unpacking and so on.) I haven't thought enough about it to go into any further detail now; but it looks as though 1000 bits per second would be a nice safe guess for the text, and the earlier argument suggests that the basic control system load would fit into that very comfortably. If we want to preserve the four-byte packets, we have to double that rate.

From the network's point of view, there is no special problem in coping with a **mobile control terminal**. The function of the terminal is to convey instructions from the person controlling the system to the various units which comprise the system, and to receive their responses. It is therefore just another unit exchanging messages with the network, so as we have allowed for any device to be mobile, terminals are included.

The rules for **initiating transactions** are "higher level" matters; provided that transactions are possible, it will be possible to administer them somehow. The main matter of immediate concern is that the administrative load on the network should not lead to a serious increase in traffic. Given the traffic estimates above, there should be room for a modest expansion without difficulty.

The requirement for **mobile devices** is met by the provision of local wireless communications. Observe again the possible need for "cellphone" operation if the mobile device is required to move from room to room. (Which it is - not to provide for such an obvious requirement would be absurd.)

THE DESIRABLE FEATURES.

There is no difficulty in principle in the network's handling **multiple terminals**. Provided that we always know the destination of a message, we can contrive to get it there. In practice, multiple terminals imply the possibility of correspondingly multiplied data rates, so there is a limit to how many terminals can be accommodated without deterioration of service. There is also the possibility of conflicting instructions being transmitted from two of the terminals; that's harder.

Ability to handle multiple terminals may also be an advantage in two other desirable activities : coping with interference from adjoining systems, and handling instructions received through the telephone.

It is far from clear that the same technique could be extended to **free-ranging garden tools**. (Unless, perhaps, they were only used at night ?) I pass.

Provision for **failure detection** is certainly desirable, and possible in principle, at least to some degree. There are two possible approaches : the initiative may be taken by the device, or by the control system.

- Fault reports issued by the device are clearly limited by what the devices can do; if the device doesn't work, it isn't going to issue anything at all. Accepting that, though, a comparatively clever device could return fault messages to the network controller through the system as already envisaged. A C- or stupid device, on the other hand, won't do anything, so unless some sort of sensor can be incorporated into the controlled on-off switch, we aren't going to get much information. Intelligent on-off switches can be imagined : it would generally be possible to detect short circuits and open circuits, and more elaborate units might be able to recognise deviations from standard patterns of operation - such as a thermostat which never turned off. All this is making the units more complex and more expensive, though, and one must ask whether it's worth while. There are also questions of practicality : the network mustn't jam with fault reports just because someone unplugs an electric heater for a while, though perhaps a short circuit report should be treated as urgent.

- The alternative is for the system to search for faults, presumably by polling the devices at some convenient rate - perhaps once a minute is adequate ? With the estimate of 80 devices used in the preceding argument, that gives an additional network load of about three messages (the poll and the response) per second. It is interesting that this rate is comparable with the maximum rates envisaged in the earlier traffic estimate. This polling system could usefully be run as a general status check; each unit polled may respond (if at all) with status information. The open circuit information considered in the previous paragraph could be included here, and it is then up to the controller to decide on the significance of the report.

IMPLICATIONS : WHAT ELSE MUST BE DONE ?

I've discussed the requirements of the network and something about the terminal devices, both controlling and controlled. I haven't said anything about **how the network should be administered**, or by what.

The basic network itself consists of a conductor, which happens to be the mains wiring, linking several units. Some of these units are infrared transducers, which must transmit any message arriving along the wire, and copy any infrared message received onto the wire. As any device may initiate a message, there is a possibility of clashes, which must somehow be resolved.

I shall take it that the network's job is to move packets of more or less the sort I suggested (four bytes, with an address in the first) from source to destination. There are many ways to manage this. Some, like ethernet, require only a communications medium to link the communicating units, and need no explicit medium controller - or, more precisely, the medium controller is distributed in the units, which must therefore be comparatively intelligent. For our purposes, that isn't a very good idea, because it is bound to increase the complexity and cost of each of the terminal units.

It seems more likely, therefore, that an identifiable network controller will be necessary. There are other reasons for making this choice :

- Something must oversee the network status polling operation, and be able to interpret the results obtained (or their absence) in a fairly intelligent way.
- Something must handle clashes.

- Something must be in a position to decide between conflicting instructions, should they arrive simultaneously from different units.
- Similarly, something must identify identical instructions received by two or more infrared receiver, and accept one of them.

I don't intend to go on here to discuss any details of how these miracles are to be accomplished, but it's clear that some reasonably capable system will be needed, and that it isn't something which one can cheaply distribute round the several terminal units in use.

If we have a network supervisor, which is presumably a computer of some sort, **where shall we put it ?** An obvious possibility is to combine it with the control terminal - but on further inspection this is seen to be by no means necessarily a good idea.

- We want the control terminal to be light, and cheap. Building an ambitious controller into it isn't the best way to achieve the desired end.

- The control terminal is on the end of an infrared link. It is likely to be safer to have the controller as firmly attached to the network as possible - which is to say, attached directly to the mains wiring.

- There is no reason to suppose that the person who normally uses the control terminal will be knowledgeable about the computer system. That being so, it's likely to be safer to put the network controller out of the way.

We've made some provision for failure of the various network units, and we expect that any such condition will become known to the controller quite soon after it happens. But what can we do in case of **failure of the network controller** ? The answer depends on the level of service we require to be maintained if the controller fails.

- If we don't mind the whole network becoming paralysed, we can manage without any special provision. This is cheap, but obviously not a satisfactory solution for anybody, and particularly so for someone with disabilities who must rely on the computer system for daily living.

- If we can manage with a reduced level of service, then a simple backup computer may be useful. This processor would poll the main controller from time to time, and take over some of its functions if it failed to respond. Perhaps it would not be able to support the comparatively high data rates needed by the more demanding activities, but it could well keep the basic environment control functions running. The additional intelligence of a second processor could also be used to call for help if the main controller fails, if that's possible.

- If full services must be maintained, then the main controller must be duplicated, and hardware or software provision made for the two units to operate in parallel and to check each other from time to time.

What are the special implications of our desire for **an extensible system** ? Our adoption of device addresses rather than fixed wiring helps; providing for the system to allocate these addresses on demand when adding a new device would be even cleverer, but would require some intelligence in the device unit, which might make it more expensive. Giving each device an address which can be set by switches may be a good compromise.

Cellphonery will be necessary if mobile devices, which includes the control terminal, move from room to room. When such a device is visible from two or more infrared transducers, its messages will appear as contenders for the network unless there is some means for determing that they are identical. This may be a job for the network controller.

IMPLICATIONS : HIGHER LEVEL STRUCTURES.

In the REVIEW I mentioned three **higher level properties** : the need for an extensible system, the communications protocol, and the need to use any sort of device. I do not propose to design them here and now, or at all, but here are some thoughts.

We will have an **extensible system** if - subject to physical limitations such as channel capacity - we can attach a new device with the minimum of change to the existing system. To a great degree, the choice of a communications bus forces us to develop an extensible system, as we can't rely on physical conductors to identify devices, so we must use addresses of some sort. Once our communications include addresses, our system is extensible at least within the confines of the number of bits we've allowed for the address. How many distinct devices should we expect the system to handle ? Without carrying out any analysis but guesswork, I don't see me wanting to drive anywhere near 50 devices. If the network is also carrying information for autonomous devices which don't need my control, perhaps we could add another 50 to feel reasonably safe. A seven-bit address will cover that; it seems a fairly safe guess that by the time we exhaust the capacity of a one-byte address it will be high time to redesign the network anyway.

Allocating the addresses is another topic. While there's nothing difficult about it, a method which can be used by unskilled people would be best. It's easy enough to imagine a variety of ways in which a device could allocate its own address; it's less easy to see how this could be done without complicating the system and thereby adding to the expense.

The **communications protocol** must be designed to cope with messages of many types exchanged between a network controller, a control terminal (perhaps more than one ?), and several devices, some of which may respond to requests for information and some of which may initiate transactions spontaneously. Matters to be settled include message types, message format, transaction structure, and how to avoid or resolve clashes.

- Message types : In deriving estimates of network load I invented a four-byte packet, but that was only a device to yield some plausible numbers, and it's unlikely to be all we need. The obvious exception is the text string. If we want to use four-byte packets for that, we quadruple the required data rate.

- Message format : The messages will have to include an address, and probably a message type. After that, the details are up to the devices concerned. Should we lay down standard message formats from the start, or let patterns evolve as required ? Should every message include error detecting information ? If so, how, and what do we do about it ?

- Transaction structure : Do we want acknowledgments for messages ? What about requesting retransmission of faulty messages ? Do transactions initiated by devices have the same structure as those initiated by the control terminal ?

- Clashes : In any network in which a passive medium can receive messages from more than one potential message source, clashes are possible in principle. We can either impose disciplines to make sure that they don't happen, or we can find ways to deal with clashes when they occur. Multiplexing methods come into the first category, but are unlikely to be practicable with this system. We have only one physical carrier, which is not well suited to frequency multiplexing, and time multiplexing by anything like a polling or token-passing method requires comparatively complex terminals and perhaps a lot of time. Equally, though, conflict resolution requires that the terminals be able to detect conflict and do intelligent things about it. I don't know that there's a simple answer.

  There may be a mildly complicated answer. It may be possible to classify the terminals according to their communications needs as well as their intelligence; then, for example, an intelligent terminal which wishes to initiate a transaction must listen to the line traffic and choose a sensible time to do it. I haven't worked it out. It seems likely that the network controller must play a part in handling clashes, but I don't know what it is.

I've already touched on the requirement that **any reasonable device** should be attachable to the system in the first section (REQUIREMENTS) of this note under the heading "What". There I identified the obvious primitive requirement : a controllable mains on-off switch, useful for the stupid machines. For the C- machines, the button-pushing dial-reading etc. robot is a serious suggestion, but not part of this project unless we have a lot of spare time. (And money.) C- machines with controllers may be more accessible; it may be easier to subvert the controllers electrically than to press buttons, though it isn't easy to see just how to devise an approach which is useful for a wide range of different sorts of controller - other than pushing buttons ! Intelligent machines will have interfaces, which we can reasonably assume will be more or less conventional serial or parallel channels. For these, we could use a rather cleverer standard box equipped with standard interfaces which could relay signals in either direction as required. It would almost certainly need to be programmable to some degree to cater for different data rates, different encodings, buffering if necessary, and so on.

RATHER TENTATIVE SUMMARY.

So far I've discussed bits from the bottom three (?) layers of the OSI network structure. Every time I think of anything at a higher level, it has some effect on the lower level bits. That isn't what I intended : I began with reasonably high-level specifications, but then dropped down to the bottom levels and bounced up again. For example, how do you get from the SPECIFICATION to the conclusion that you really need a network controller ? Is there a systematic way of doing it ? (Probably there is, but these things don't come my way.)

Lots of stuff about the network management remains undefined. Just what is the function of the network controller ? In some sense it's where the network's intelligence lives, but what sort of intelligence do we need ? It has to take as much load as we can manage from the terminals - particularly the control terminal - so that they can be as simple and cheap as possible, but I doubt whether any more progress in that direction is possible without a much clearer idea of what we want the system as a whole to do.

Most important : Despite the many unresolved questions, I think it looks possible.

REFERENCES.

1 : C. Walker : "MARC mains appliance control system", *Everyday Electronics* **19**, 373, 457, 526 (1990)

2 : C. Walker : "MARC phone-in", *Everyday Electronics* **20**, 84 (1991).

3 : M.A. Regalbuto, J.B. Cheatham, T.A. Krouskop : "A framework for a practical mobile robotic aid for the severely physically disabled", in reference (7), page 127.

4 : G.A. Creak : *A view of rehabilitation computing*, Auckland Computer Science Report #46, Auckland University Computer Science Department, 1990.

5 : M.M. Trivedi, C-X. Chen, S.B. Marapane : "A vision system for robotic inspection and manipulation", *IEEE Computer* **22#6**, 91 (June 1989).

6 : J. Schauer, D.P. Kelso, G.C. Vanderheiden : "Development of a serial auxiliary control interface for powered wheelchairs", in reference (7), page 191.

7 : *RESNA'90 : Capitalizing on technology*, Proceedings of the 13th Annual Conference of the Rehabilitation Engineering Society of North America, June 1990.