

DOES CENTRALISED CONTROL MEAN CENTRALISED COMMUNICATIONS?

or

DID MURPHY ERR?

by John Aitken BE, MIEEE, AMIRSE
Director
Aitken & Partners
Consulting Engineers

IRSE Technical Meeting, Sydney, 27 April 1990

Prologue

The Sydney North Shore Times recently related a tale which should bring a chill to the spine of communications and signals engineers. I quote:

Police say the series of events began at 3.00 pm on April 9 when the man visited a patient at Royal North Shore Hospital, St Leonards.

As he was leaving, he saw a Chubb safe in one of the hospital offices.

He placed the safe and a key locker on a hospital trolley and towed them to his car.

He tied a rope around the safe and the trolley and tied them to the car bumper. He then towed the trolley along the Pacific Highway until the wheels fell off the trolley.

The man got out of the car, unhooked the trolley and tied the rope to the safe, dragging it another 400 metres as sparks flew off the roadway.

When he reached his home unit, he hired an angle grinder and cut the backing plate off the safe.

Lane Cove police went to the home unit after the man's neighbours had complained about the noise.

They found the man in a distraught state who told them he had found only \$20 in the safe after he had finally opened it.

Police said it cost the man \$50 to hire the angle grinder.

Communications and signals engineers are known and respected for their thoroughness and careful planning. They do not act on impulse but plan and design systems which are robust and have "fail-safe" operation and redundancy built in.

None the less, the unforeseen does occur. It may pass un-noticed, it may be identified as an incident or it may become a disaster. The impact is nearly always determined by the timing. If confirmation of this is needed one need only refer to Mr Nock's book "Historic Railway Disasters" or any issue of the "Aviation Safety Digest".

Centralised Control

Centralised Control is a favoured approach to railway operation. Spawned by the technological advances in communications systems, centralised control has seen

- improved utilisation of the track
- improved utilisation of rolling stock
- reduced staff levels.

This may be summarised as more efficient operation with lower costs. On the other side of the ledger we see the disadvantages

- reduced headway between trains
- less staff available in the field
- greater reliance on communications.

I consider these factors to be disadvantageous in the present context because they are factors which can combine with the unforeseen to result in disaster. Few railways have replaced signal box and station staff with hot box detectors and dragging equipment detectors. The "eyes and ears" which were once located at frequent intervals along the track have gone and, in many cases, have not been replaced by alternative means of surveillance. When there is a fault or problem it may be quite some time before it is noticed and disastrous consequences may have occurred.


The consequence of all this is that centralised control has brought about a very substantial reliance on communications and especially on radio communications.

Centralised Communications


Centralised control systems are always designed for reliability and some degree of protection from any single failure. The topology of the network will be some hybrid of the common schemes depicted overleaf.

The particular combination of the schemes used will usually be determined to some extent by


- the economics
- diversity required
- reliability required
- reliability of individual paths
- future plans.



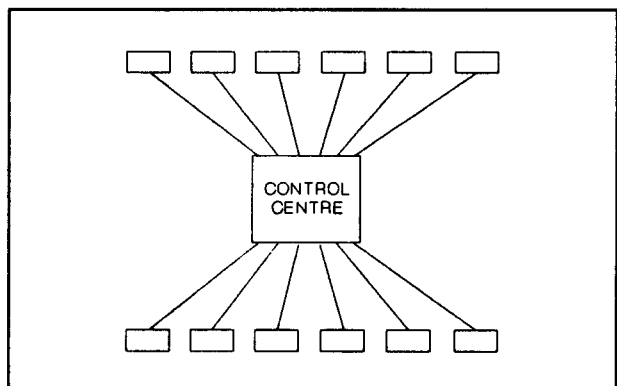
Parallel Feed



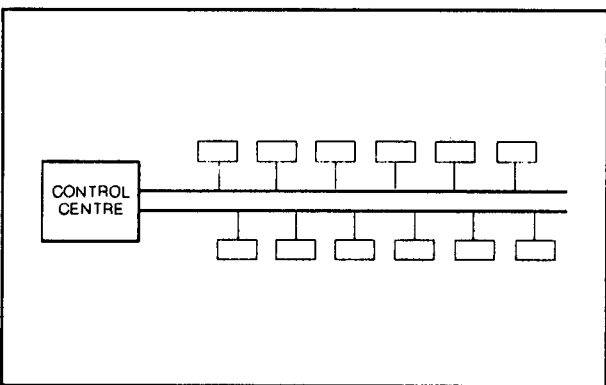
Parallel Feed in Sections




Ring Feed



Star



Alternate Feed



Star + PSTN

Each of these schemes has at least one weak point. The Control Centre.

Engineers will plan for diverse cable routes, redundant power supplies, multiple processors, independent software packages, alternate ventilation systems.....all manner of techniques to protect the centralised installation.

Despite all this the whole system is vulnerable to disruption and destruction because the absolute protection of a single site is beyond the control of the design engineer.

Consider the threats to which a site is exposed. An entire control centre can be reduced to inactivity or uselessness by

- fire
- smoke or toxic gas
- earthquake
- bomb threats
- malicious damage
- malicious EMI
- lightning
- flooding.

While the majority of my audience is now relaxing because they know that

their control centre is protected against fire
it has no smoke-producing cables
earthquakes never happen
no-one bombs railways
graffiti does not stop trains
no-one produces malicious EMI
lightning does no damage and
the centre is on the top floor;

the major banks are all building or have established standby and disaster-recovery computer centres.

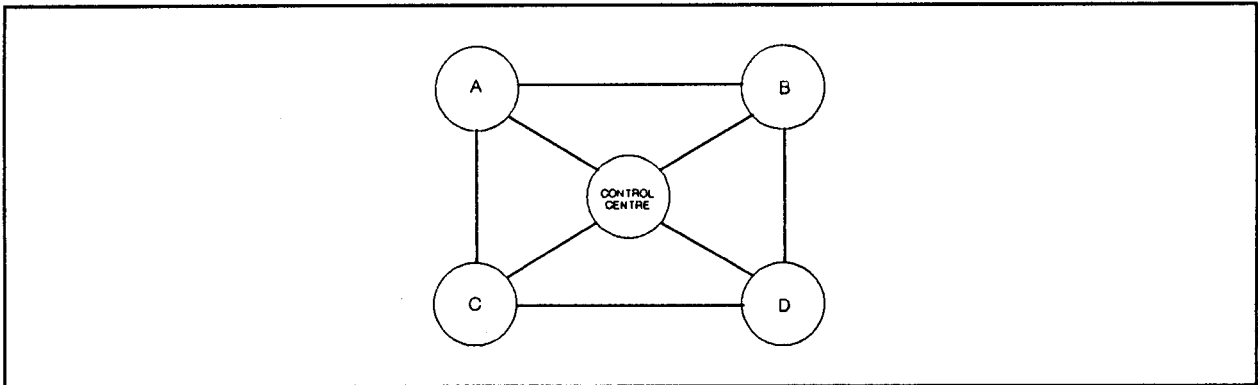
The banks are no less interested in reliability than a signalling engineer but the outlook is different - a stopped railway is safe; a stopped bank is broke.

The increasing pressure for reliable on-time performance in freight delivery may well change the outlook of the railways.

Let's return for a moment to the disadvantages which have come from the introduction of centralised control. Conventional centralised control schemes have local panels at interlockings so that operation can be maintained if the overall system has failed. It is not sufficient to have local control at interlockings if the staff to operate these have all been given golden hand shakes and the train has driver-only operation. Communications are vital and an approach is needed which allows normal operations to be resumed quickly after the loss of the control centre.

If it is not practical to have duplicated control centres, and in many cases it may be, then we should at least be striving to provide decentralised communications so that at least some degree of control can be quickly re-established after a disaster.

While this seems obvious, centralised control has an insidious influence on communications systems. The PABX network will have a node at the control centre. This is, of course, logical and economic. This node will connect to the other nodes in the system with a common signalling system so that networked operation is achieved. With a common numbering scheme the result will be a network which is highly adaptable and robust.



Nodal Network

If the design is one which employs a ring approach as well as a hub then many desirable features have been achieved. All that is needed now is the ability to re-configure the network so that say, Node A can become the control centre and perhaps disaster recovery can be achieved.

Perhaps it can be achieved: the signalling control and communications, the data files the processing all need to be associated with Node A. This requires a lot of will and financial conviction because a duplicated control centre is not just a matter of skilful network design. Or is it?

What can go wrong now?

The scheme just described appears to solve the problem. We have duplicated centres in diverse sites with diverse, multiply-fed communications nodes and a simple changeover mechanism. What could go wrong?


What could go wrong now is that the network could fail. It could fail because

- independent routes are not independent
- the programming has not been kept up to date
- the network has developed a software fault.

The first two problems have self-evident solutions. But there is a moral. The network concept needs to have regular maintenance just as much as the network components. Only a disciplined review of all aspects of the network design at regular intervals will ensure that the network retains its original diversity and independence.

The third problem is harder. Increasing network complexity is introducing new hazards as the control systems have more scope for bugs. Two approaches are noted

- duplication of networks
- "dismantleable" networks.



Duplicated Nodal Networks

Duplicating the networks is an approach which is relatively expensive but has advantages. The concern must always be that the duplicated networks are truly independent and remain that way. The quick solution to a fault or a temporary overload at one node, which uses a part of the duplicate system, may be acceptable in the short term - in the long term it may render the independent system very dependent.

The other approach is to find a means of dismantling the network so that it can be re-set and then re-constructed. This is no easy task but is something which must be considered.

Does centralised control mean centralised communications?

Did Murphy Err?

Epilogue

I conclude with another newspaper article. This time from "The Institute", the news supplement to IEEE Spectrum.

Vulnerability Exposed in AT&T's 9-hour Glitch

A bug in a new version of switching software generated AT&T Co.'s most extensive interruption in service ever, reawakening concerns about the vulnerability of today's highly centralised US phone system.

The scale of the failure, which was sparked by a trivial hardware problem, surprised experts and the public alike. The AT&T network could handle only half the normal volume of long-distance, international, and toll-free 1 - 800 calls in a nine-hour period beginning at 2:21 pm on January 15, the Martin Luther King birthday holiday.

The problem is a consequence of the concentration of network control in fewer points around the country and the increasing reliance of phone network switching on software for signalling. Both allow sophisticated calling features and more efficient operation at lower cost.

"As the phone companies move in the direction of centralising decision

making to optimise usage of network, it becomes easier for this kind of thing to happen," said Amos Joel, a retired AT&T Bell Laboratories engineer who pioneered modern switching techniques.

Switch Glitch

The instigator of the breakdown was a minor fault in a call-switching system in lower Manhattan that triggered a software bug at other switching nodes around the country, so that the problem cascaded throughout the network. The processor that controls the operation of the New York City switch noted a hardware problem in the interface between the switch and the signalling network, according to Karl Martersteck, vice president of network development for both AT&T and AT&T Bell Laboratories.

In a routine procedure, the processor then took the switch out of the network temporarily while it investigated the problem, sending messages to other switch processors about its shutdown so they would route calls around it. All this command and control signalling activity was on a separate channel from the voice and data information, an innovation called common channel signalling that AT&T began installing in 1976 to speed operation.

The New York switch came back online within seconds, and its processor sent messages to other switch processors on the signalling network announcing the arrival of calls. Normally those processors would then change the status bits for the New York switch back to "online." AT&T investigators deduced that a processor at a connecting switch was still changing status bits when it received a message announcing another call (within 0.01 second of the first message). Because of a flaw in software logic, the processor became confused and shut down so it could reinitialise.

When that processor came back online, the error repeated on another switch and then another, spreading rapidly throughout the network. It was triggered whenever the first two calls coming from a downed switch occurred within 0.01 second of each other. Martersteck said the probability of such rapid-fire calls occurring at so many switches was extremely low, even during the Monday afternoon peak calling period. As the failures spread, the number of control messages "flying around" the signalling network mounted and exacerbated the congestion, he added.

The 4ESS switches, the principal switches in the AT&T network, are each designed to handle 700,000 calls per hour, which works out to 0.005-second intervals, so the software should not have balked at intervals twice as long. Ironically, the new software was installed in December to increase the network's reliability by allowing auxiliary signalling channels to be used when others are congested or down.

Engineers at AT&T's network operations centre in Bedminster, NJ, and technical support centre in Lisle, Ill, became aware of the problem immediately, Martersteck said, when remote monitoring equipment lit up

congested lines in red on a large computerised national map display. They followed standard procedures for correcting problems, but for the first time in AT&T's history, none succeeded.

Working in teams, specialists from Bedminister, Naperville, Ill, and Columbus, Ohio, analysed the data and tried to pinpoint the source of the problem by systematically isolating portions of the network and initialising individual elements within the segment. "We had to be careful we didn't lose the ability to handle calls and [to] make things worse," Martersteck said.

By 4:30 pm, the teams had determined that the problem was located in a signalling network, and they then began reconfiguring it for stability. They cut off the excessive flow of control messages by disabling selected commands and by disconnecting network signalling links to physically block the messages, Martersteck said. The next day AT&T replaced the new software program with the old version and on January 20 loaded new software for a permanent repair.

Sophisticated Sorrows

Though AT&T's latest technological innovations have undeniably improved phone service, experts say the network's sophistication and complexity must take some of the blame for the confusion the software bug caused. The software associated with one 4ESS switch consists of more than 2 million instructions, according to Joel, and there are now 114 of those switches on AT&T's network.

Joel said the interaction between the signalling network, known as SS7, and routing control centre is very complex and, consequently, unpredictable. "Routing control is the only thing that knows what trunks are being used," Joel said.