

IEN: 139

J. Haverty
BBN
April 1980

HOSTs as IMPs

Jack Haverty

Bolt Beranek and Newman Inc.

The following discussion is an extract from a report for one of the ARPA projects at BBN. The relevant details of the system configuration include the following:

- o a PDP-11 with IMP11A interface is used to communicate with an LSI-11 using a Collins "1822 interface"

- o the electrical connection uses the "distant-host" configuration

This report was originally issued in October 1979. It is being reissued now as an IEN after some recent inquiries which indicate the problem may be more widespread than was previously thought.

April 1980

HISTORY

We've been having an enormous amount of trouble (more than usual) since early summer getting the LSI-11 and PDP-11 to talk reliably on the 1822 lines. When the problems were investigated, the IMPTST program also reported errors on the 1822 connections. This has happened in the past, and often the problem would disappear after swapping some cables, or maybe waiting a few hours. However, this time the problem was persistent, which at least gave us the opportunity to figure out what was happening.

A collection of people have been working on the problem, and we've come up with an analysis of the situation.

The immediate cause of the strange behavior was tracked down to large noise spikes induced in the cables. We have noise spikes every half-cycle on the building power cables. These are coupled into the 1822 cables by RF coupling. We have seen as much as 20 volts peak-to-peak spikes on the differential signal lines. This appears to the receivers as common-mode noise.

Differential receivers in the various host interfaces have receivers which are spec'ed for good common-mode rejection up to some maximum input signal level. For the LSI-11s, the chips are good up to

+15 volts. IMP11As use chips which are spec'd at +-3 volts.

For the various "IMP" interfaces, the PTIP interface is good for a range similar to that of the BCRs good for a similar range. The 316/516 IMPs provide optical isolation, so they are good up to some very high value dependent on the breakdown of the insulation.

The 1822 spec says that host's receivers should expect a 1 volt peak-to-peak signal, centered around the host's ground. Host receivers also must be able to tolerate at least 2.5 volts of common-mode noise. Note that, for the IMP11As, DEC picked chips which exactly met the spec, i.e., 2.5 volts noise plus 0.5 volts signal in either the plus or minus direction requires an input range of +-3.0 volts. The chips in that design have a maximum legal input range of +-3 volts with respect to the PDP-11 ground. Thus the DEC design is a legal host interface according to 1822.

By judicious grounding and other magic, we have gotten the noise down to less than a volt. The source was traced to our building UPS system, which has some bad filter capacitors. These will be replaced, and that is likely to remove the noise source. However, 1 volt of induced noise is not unreasonable, and is also probably fairly common in computer areas. The UPS problem only makes the noise occur every

8 milliseconds, whereas more traditional noise might occur at intervals of seconds or minutes.

As of April 1980, the filter capacitors have been replaced. This has reduced the noise somewhat. We have learned however that UPS systems can generate significant noise in normal operation. This noise appears on the input side of the UPS system; the output side is generally well filtered. Computer equipment powered from the UPS output will not see any significant noise on its power; equipment powered from non-UPS power will have the noise present, but the equipment itself typically filters the power enough so that the logic power is noise free.

The problem we have been investigating appears to result from coupling between "1822" cables and power cables carrying non-UPS power which also carries the noise spikes generated by the UPS system.

HOSTS AND IMPs

With the 1 volt noise, the PDP-11 and LSI-11 communication still doesn't work properly. IMPTEST reveals occasional errors.

In tracking these problems down, we delved pretty deeply into all the pieces of the system, even down to the circuits used within the

driver/receiver chips. In the IMP11A receiver, it appears that if the signal level exceeds 3 volts, the logic will get "confused", and produce an inverted output. However, with 1 volt of noise, and 0.5 volts of signal, the input should not exceed about 1.5 volts.

Further tracking along, we looked at the drivers in the LSI-11. These are single-ended drivers, i.e., using a single +5 power supply. They produce signals which swing between +0.2 and +2.6 volts. The upper limit will depend on the actual value of "+5", resistor tolerances, etc. The measured "differential" signal produced by a pair of these drivers keeps one signal line at +0.2, and the other at +2.4 all the time. Thus the signal is a 2.2 volt peak-to-peak signal, with a built-in common-mode noise level of 1.2 volts.

When an LSI-11 is connected to an IMP11A, the IMP11A receivers see a signal which swings to +2.6 volts. Given 0.4 volts of noise, this drives the receivers beyond their specifications.

We believe that this is the reason behind the current problems we are having with the link. With the ambient 1 volt noise, the IMP11As receivers are being driven out of spec all the time, which causes random errors depending on how the signals add, other noise,

etc.

We have also experienced in the last few years a continuous problem of random incompatibilities between the LSI-11s and the IMP11A. For example, one day a particular LSI-11 might work well to the PDP-11, and the following day it might not. It seems conceivable that this is partially caused by the situation I just outlined. The real voltage level on any day might depend on the building power voltage, temperature, noise produced by other equipment in the vicinity, etc. Since the IMP11A is being driven at the very edge of its working range, all the random "noise" factors will determine how well the link works at any time. We have observed also that the system occasionally "hangs" for no apparent reason, possibly after running for hours, which may also be related to the aperiodic noise which might occur as air conditioners switch, elevators, etc. A similar installation has also reported problems of this nature, with unreliable LSI-11/PDP-11 communications using distant-host connections.

CONCLUSION

The consensus of opinion here is that the problem lies with the LSI-11 1822 interface characteristics. The "1822 spec" dictates how to build HOST interfaces, but does not specify how to build IMP

interfaces. It alludes to the fact that the IMP provides for things like ground isolation, deskewing, and other features not required of hosts.

The 1822 interface is however asymmetric -- it is NOT true that all host 1822 interfaces should be able to talk to all other hosts' interfaces. The host spec defines an interface which will allow any host to talk to an IMP; the IMP interface, which has to talk to ALL POSSIBLE host interfaces, performs more functions than a host interface is required to do. In particular, it is required to provide ground isolation, and to generate true differential signals, centered on the host interface's ground which is carried by the ground-wire in the cable.

The LSI-11 1822 design seems to adequately follow the 1822 spec, in that it provides a host interface. The IMP11A also adequately follows the 1822 spec, as a host interface.

However, the LSI-11 must act as an IMP -- i.e., it must be designed to communicate with any "1822 interface" which meets the host specs. The current LSI-11 1822 design does not meet this criterion.

The root of the problem is probably that most people think of the 1822 specification as a symmetric one, i.e., such that any 1822

interface can talk to any other 1822 interface. For most interface implementations and environments, this is probably true, which reinforces the mistake.

SOLUTION(s)?

We are looking at various short-term solutions to see which is the least painful way to make the system function reliably. In the long run, we believe that the LSI-11 1822 has to be changed to behave as an IMP. At the very least, the drivers must generate a ± 0.5 volt signal, centered around ground, and provide for ground isolation. There may be other constraints, such as timing issues, as well, which an IMP designer could help with. We have not looked into the LSI-11 1822 design in these areas.

It is also worth noting that other projects which use "pseudo-IMPs" should be examined, to see if the same situation exists. In the LSI-11-IMP11A case, the fact that the system is on the edge of the working range means that it works most of the time, and is just annoyingly flaky. Other people may have the same situation.