

Network Working Group
Request for Comments: 1216

Poorer Richard
Almanac Institute
Prof. Kynikos
Miskatonic University
1 April 1991

Gigabit Network Economics and Paradigm Shifts

Status of this Memo

This memo proposes a new standard paradigm for the Internet Activities Board (IAB) standardization track. Distribution of this memo is unlimited.

1. Introduction

The history of computer communication contains many examples of efforts to align the capabilities of processors to that of communication media. Packet switching is the classic case of a careful tradeoff between the costs of memory, processing, and communications bandwidth.

With all of the attention and publicity focused on gigabit networks, not much notice has been given to small and largely unfunded research efforts which are studying innovative approaches for dealing with technical issues within the constraints of economic science. This memo defines one such paradigm.

2. Contemporary Network Economics

Recent cost estimates predict a continuing decline in the cost for processing, memory, and communication. One recent projection put the decline for \$/bit and \$/MIP at 99% per decade and put the decline for \$/bps at 90% per decade. Scalable parallel processor designs may accelerate the cost declines for CPU and memory, but no similar accelerated decline should be expected in the cost of communications. Such a decline would imply eventual declines in the cost of 56Kbps service used for voice, resulting in a negative rate of return for telecommunications carriers, an unlikely eventuality even if free-market forces are carried to their logical extreme.

Increases in processing power create additional demand for communications bandwidth, but do nothing to pay for it. While we will sell no paradigm before its time, the 9% difference, particularly after compounding is taken into account, will bankrupt the internet community unless a paradigm shift takes place.

3. The ULS Paradigm Shift

The ULS paradigm shift breaks the downward spiral by concentrating on end-to-end datagrams and virtual circuit services operating in the .01 uGbps region, namely Ultra Low Speed networking.

However,

"The worlds best technological paradigm shifts are useless unless they (a) are economically viable, (b) have clear applicability, (c) are technically feasible."

--Milton John in "Paradigms Lost"

3.1 Economic Viability

Cost projections indicate that individual ULS circuits can be provided at a cost of <\$.03/month due to the unusually high multiplexing that will be possible on Gbit links. The 10 THz bandwidth of existing optical fibers will be able to support on the order of 1 TUser, handling population growth, and even internet growth, for some time. Moreover, if \$.03/month is a significant barrier to entry, substantial discounts appear to be economically feasible.

3.2 Clear Applicability

A fundamental principle of networking is that network speed must match the application. We have identified a number of critical applications that are matched to ULS technology. Below we itemize a few of these, but we provide a brief description for only the first; the match for the others should be equally obvious.

- Low priority facsimile: A large percentage of documents and letters are sent via facsimile not because they need sub-minute delivery, but because they carry signatures or graphics. In these cases, a three-hour delivery (comparable to the value reliably achieved on many of today's packet-based email systems) is sufficient. With proper compression, this delivery time can be achieved over a ULSnet.
- Real time data (e.g., tracking glaciers)
- US postal service
- Contracting for research

To be truly viable, ULS networking must scale, and indeed it does.

With some effort, we envision extending the technology to the extremely-low-speed regime. Applications that scale from the ULS applications above are:

- Real time data (e.g., gravity wave detectors)
- Italian postal service
- Congressional budget process

3.3 Technical Feasibility

The hardware issues are well in hand. The remaining issues are protocol related. To examine them, we must extrapolate backward from some well known networking principles.

"Gigabit networks require new protocols."

The clear inference here is that ULS will require old protocols, so as we recede into the future, we should expect the following:

ULS will require minimal development. Although we may need research in storage technology to recover the software from old media such as decayed magnetic dump tapes, paper tape, and partially recycled card decks, this effort will be more than offset by the savings.

ULS protocols will be well documented, amenable to verification, and suitable for MSI implementation in Silicon, or even Germanium or relays. In particular, the alternating bit protocol [1] is a leading contender.

"Bad news travel fast."

Therefore, ULS gives preferential treatment to good news. While this will delay the delivery of bills, notices from timeshare condominiums, and contest announcements, it will also produce immediate productivity gains on several mailing lists.

3.4 Problems Requiring Work

ULS is not without problems.

Some other well-known protocol suites are well ahead of ULS in exploring the desired performance operating point. We note our concern about the dearth of domestic (U.S.-based) research and development in this important area. This is particularly disturbing in light of the level of work now underway in other countries.

Efficiency is a problem:

- All ULS protocols incorporate slow-start.
- Lower data rates mean fewer errors.
- Whereas modern protocols use 32 bit sequence numbers, acknowledgment fields, etc., ULS headers can be quite small (1 bit sequence numbers for the alternating-bit protocol). Thus the header/data ratio shrinks.

The net result is "creeping efficiency" which tends to push us away from the proper ULS operating point. While we have no definitive solution, there are several promising palliatives:

- Forward Error Insertion (FEI)
- Negative window scaling factors
- New protocol layers
- Multiple presentation layers

4. Conclusions

The road to Ultra Low Speed (ULS) technology is long, slow, and easy.

REFERENCES and BIBLIOGRAPHY

- [1] Lynch, W. "Reliable full-duplex file transmission over half-duplex telephone lines", CACM, pp. 407-410, June 1968.

Security Considerations

Security issues are not discussed in this memo.

Authors' Addresses

Dr. Poorer Richard
Almanac Institute
Center against Misoneoism
Campo Imperatore, Italy
EMail: none

Prof. Kynikos
Miskatonic University
Arkham, MA.
Email: Kynikos@Cthulu.Miskatonic.EDU