(draft)

Some Thoughts on Utilisation of QUESTnet Links

There has been some concern expressed as to the justification of 2 Mbps links on Questnet,.People have pointed out that the traffic on these links does not approach the 17 Gigabytes per day possible using such a link. The quite legitimate concern has been expressed that we will have "an enormous network with no services on it.".

While there certainly are grounds for concern here, I believe it can be shown that 2 Mbps bandwidth is already justified for all institutions with the possible exception of UCQ and USQ, and that even here 2 Mbps will soon be justified.

In this document I present a few of the criticisms that can be levelled at the dimensioning of QUESTNet with the answers to these criticisms.

"2 Mbps is Overdimensioned"

If one performs traffic measurements on QUESTNet links, one typically finds that during the entire busy hour, is about 2-5% of a 2 Mbps link. The naïve reaction here is to say that we could survive with 100 kbps links instead of 2 Mbps links.

This analysis, however, ignores a number of important issues.

(a) Difficulty of Aggregation

Bandwidth, whether bought from Telecom or provided using microwave links comes in standard chunks. One can provide 48 kbps, 64kbps, and 2 Mbps channels. It is possible to provide channels of size between 64kbps and 2Mbps using channel aggregaion equipment to aggregate 64kbps channels. This solution, however, has several serious technical advantages, and the equipment required costs in excess of $15 000 per site.

Ignoring the technical disadvantages of aggregated channels, Megalink becomes cheaper than multiple ISDN channels at a bandwidth of about 900kbps (depending on distance). For links within the metropolitan area, Megalink becomes cheaper than 128kbps.

(b) Delay Characteristics

There are many figures of merit that can beused when describing packet-switched networks. The history of academic networking in Australia has stressed the throughput metric, since the networks used by most Universities only 3 years ago were only capable of file transfer, electronic mail, and closely related services. Services where delay is critical such as remote terminal access were not provided.

The installation of QUESTNet has allowed services hitherto infeasible to be offered, and end-end delay has become an important figure of merit. The formula used by teletraffic engineers to compute the delay for a single link (assuming Poisson traffic and constant transmission delay) is as follows:

$$W = \frac{p}{2u\,(1-p)}$$

Transmission delay is given by:

$$t = \frac{L}{s}$$

Where p is the long-term link utilisation, W is the time spent waiting in the transmission queue, and u is the rate of service (in this case the reciprocal of the transmission delay.) L is the length of a paclet, in bits, and s is the speed of the channel in bits per second. The table below shows the average packet delay (the sum of transmission delay and queuing delay) for a variety of speeds, loads, and packet sizes. All delay times are expressed in milliseconds

	Packet Length = 64 Bytes			Packet Length = 560 bytes		
Load	64 kbps	128kbps	2Mbps	64kbps	128kbps	2Mbps
10kbps	8.5	4.10	0.25	74.28	35.62	2.19
20kbps	9.58	4.26	0.25	83.89	37.34	2.20
30kbps	11.25	4.5	0.25	83.89	37.34	2.20
40kbps	14.32	4.79	0.25	125.32	41.94	2.21
50kbps	21.76	5.16	0.25	190.42	45.13	2.22
60kbps	66.40	5.62	0.25	581.05	49.25	2.22
80kbps		7.16	0.26		62.66	2.23
100k		10.88	0.26		95.21	2.25
120k		33.20	0.26		290.52	2.26
200k			0.26			2.31
500k			0.29			2.55
1M			0.375			3.28
1.5M			0.625			5.46
1.875M			2.125			18.59

The packet length of 64 bytes was chosen, as it is the size of the majority of packets sent by terminal emulators. 560 was chosen as it is the largest packet size sent by most TCP implementations, and as such is very commonly used for file transfer and screen-based interactive processing

CCITT standrds X.135 and X.136 contain objectives for packet networks, including recommendations about maximum end-end packet delay. Recommendation X.135 sets the maximu busy hour mean end-end packet delay for national networks at 350 milliseconds. This figure is barely reasonable for interactive applications. It implies a response time of the order of more than 3/4 of a second, which many users would regard as only barely tolerable. Furthermore, since packets from many sites have to traverse substantial networks before reaching QUESTNet equipment, this must be seen as very much an outside figure.

While this discussion applies to QUESTnet only, I think it is appropriate that the dimensioning of links be done such that transit delays across AARNet fall within the standard, if and when all links in AARNet are dimensioned to QUESTNet standards. This means that we must consider a packet transit comprising four hops. (Site - regional hub - national hub - regional hub - site). In the following analysis, I have assumed that regional hubs and the national hub show the same queuing characteristics as non-hub sites. The situation is a little more complex than this, but I assert that this assumption is reasonable for the purposes of this rather rough analysis. It is assumed that the packet processing time in the packet-switching nodes is negligible, which may or may not be a reasonable assumption. The upshot is that link delays have to be kept below 80 ms at the very outside.

From the table above, it can be seen that if packets are small (as in X.25 networks) this target can be achieved using 64kbps links as long as the load does not exceed 95% link utilisation. Unfortunately, many IP packets are large, and hence 64kbps links only meet the specification for loads of about 30-35% link utilisation (ca. 20kbps) or below. All QUESTNet links are running at 30% utilisation for a 64kbps channel in the busy hour. If 128kbps channels were used, a link utilision of 60-70% (ca. 80-90kbps) would provide delays which are within the CCITT recommendations. 2 Mbps links eliminate all concerns about packet delays except under saturated conditions.

(c) Burstiness of Traffic

The above analysis was made using the assumption of a Markov arrival process for packets. Unfortunately, on most LANS (whence QUESTnet receives its traffic) display traffic patterns which are considerably burstier than this. The greater the difference between the trunk speed and the LAN speed, the worse this effect is. The result is that the situation is actually slightly worse than the above analysis would suggest.

Furthermore, smaller data packet networks such as QUESTNet exhibit far greater variation in load over hourly periods than large networks such as the AUSTPAC X.25 network. Thus, the use of busy hour figures is probably not sufficient. Busy minute traffic statistics or busy ten-minute figures would probably be more appropriate. This makes the above analysis seem even more optimistic. 64 kbps links would seem to be too small for a network where considerable amounts of interactive work are being done even at loads as low as 20% link utilisation.

The results above are supported by experiences gained while AARNet links were 48 kbps. While this speed is perfectly adequate for applications like file transfer, remote logins were exceedingly painful and the use of graphical user interfaces all but impossible.

Within the next 10 months, QTAC will be "going live" with a tertiary entrance system based on X-windows graphical user interfaces. It is intended that member institutions will be able to make inquiries from their home campuses using QUESTNet. If it is QUESTNets intention to be able to support this kind of service then QUESTNet is very definitely not overdimensioned.

"There are no Services on Questnet"

Another area of concern for the QUESTNet management is the current shortage of services using QUESTNet. After all, QUESTNet was originally funded to aid the cooperation of Industry, Government, and the Research organisations. To date research organisations have benefitted, (as evidenced by the QTAC system and similar usage which is currently small, but growing fast.) but we have not seen a great deal of participation from Government and industry.

While these concerns are very valid, it would seem that the predicted participation from Government and industry is materialising slowly, with the planned connection of departments such as QDPI and the existing connections to high-tech industry representatives such as Stallion Technologies and Mincom.

It must be remember that the upgrading of QUESTNet links from 48kbps to 2Mbps has not just been an exercise in improving existing facilities. It represents the provision of completely new capabilities. Planned applications such as compressed video and audio over packet networks only became possible within the last year, and take time to happen. An analogy in public works would be upgrading Archerfield airport to the size of Eagle Farm. Superficially, both are airports, but Archerfield cannot support international traffic and it could never be the basis of a tourist industry. The facilities are simply not comparable.Add to this the fact that in implementing QUESTNet we have effectively built Eagle Farm Airport to encourage the development of an Aircraft industry, and it is little wonder that after this relatively short time it still has spare capacity.

QUESTNet is an enabling infrastructural development and as such it is unsurprising that the services that will eventually run over it are still under development.