INDRA Note 680
IEN *** 52
21st July 1978
section 2.3.3.17

Types of Service in the Catenet

C. J. Bennett

ABSTRACT: The three main traffic types each require different types of service support in the catenet. The impact of offering support for particular services on the catenet architecture is discussed. The major problem identified is supporting bulk transfer, due to the lack of gateway to gateway flow control.

INDRA Research Group
Dept of Statistics and Computer Science
University College London

Introduction

The last internet meeting raised the question of how the internet datagram protocol can support various kinds of service. At that meeting, the three standard traffic types were identified as defining the types of service we wished to support by their characteristics. The aim of this note is to identify in more detail services needed to support these traffic types, and to point out the areas of the internet and gateway protocols which will be affected by the need to support them.

Interactive Traffic.

Interactive traffic is characterised by short packets sent at irregular intervals, and requiring a response within a time determined by a user's tolerance threshold - typically, a second or so. Hence the prime constraint that must be satisfied for an interactive service is simply that the packets be delivered with the minimum delay.

The internet protocol as it currently stands does not guarantee delivery (Postel78), and the proposed form of gateway congestion control is simply gateway blocking (Cerf77). Hence, flow control is maintained by the routing algorithm. Thus supporting a minimum delay requirement means in operational terms that the traffic should be routed along minimal delay paths. This can be done by requiring the gateways to use minimum delay as the objective function of the routing algorithm. This is in fact the basis of the dynamic routing scheme suggested in (Strazisar78), and so support of low-delay traffic simply requires the catenet to operate in its normal mode, once this dynamic routing is installed.

On the other hand, the catenet cannot offer classes of minimal delay service on this basis. This needs minimisation of queueing delays rather than tranmission delays, which makes its operation very similar to priority classes. As we shall see in the next section, priority classes also play an important role in supporting stream traffic.

Stream Traffic

The second traffic type requiring service support is stream traffic. Typically, this will consist of a steady flow of data such as voice or telemetry data. The data will usually contain a high degree of redundancy, and so a degree of packet loss is tolerable. However, significant congestion leading to the loss of a large contiguous section of the data should be avoided. Often the traffic source will be of low intelligence, so techniques for retaining end-to-end reliability may not be possible in any case. Different applications of stream traffic will have different requirements on the catenet. Traffic voice, for instance, has a strong requirement for minimising end-to-end delay, though the main requirement is that the inter-arrival time of the packets should not exceed some value. Telemetry applications will normally be adding data to a large data base for later offline

processing, in which case delay is not a significant factor, although in some cases a fast response may be required. Processing efficiency will be achieved, however, if the updating process can be scheduled regularly with a high probability of there being data available to process; this applies to both telemetry and voice data.

Stream traffic, therefore, can best be serviced by ensuring that the regularity of the data flow is maintained. The role of the catenet is to avoid congestion in the gateways (and elsewhere in the catenet where possible), by giving these stream packets priority in the gateway queues. A consequence of this definition of stream traffic support is that it becomes meaningful for the catenet to offer a number of distinct priority classes, rather than simply flagging the packet as 'priority', as priority becomes a defined function of the gateways. The gateways may well attempt to support priority classes within a network in addition. In this case, the map between the standard internet priority classes and the local priority classes will be a matter for the implementation of the gateway half for the local net.

Priority is not, in itself, a guarantee that the packets will be delivered with an acceptable degree of packet loss. The gateway blocking techniques of (Cerf77) are not allied with any mechanisms for signalling other gateways, and hence it is likely that packet loss will occur in bursts. As I noted above, such bursts should be avoided; this requires more sophisticated flow control procedures than are currently used between the gateways.

Bulk Transfer Traffic

This traffic type is characterised by the movement of a continuous stream of large packets, all of which must be delivered reliably from end to end. One requirement of the catenet which has been proposed is that this stream be delivered from end to end in the shortest possible time (Cohen78). At first sight, this is the same requirement that interactive traffic has. However, there are a number of critical differences connected with the size of the transfer. Firstly, the delay constraint is much less stringent than it is for interactive traffic; the requirement is that the whole body of data be transferred in the minimum time rather than that any portion of it be so transferred. Moreover, in many circumstances the bulk transfer will be run as a batch process, in which case there may be no delay constraint at all. The potential volume of data involved also means that the capacities of the channels available do become a significant limitation on the transfer time, as they govern the point at which delays for individual packets start to increase as traffic builds up. For this reason, a throughput measure is a significant figure for bulk transfer traffic - it represents the normal operating conditions of the network.

Underlying both these comments is the fact that the service requirement for bulk transfer traffic is not packet based - it refers to the whole of the transfer. This is very different from the stream and interactive cases, where one can define service support on a per-packet

basis. Depending on the particular application we may wish to select priority and delay classes. However, the feature of the catenet that uniquely affects bulk transfer is the capacity of the networks and gateways. In a virtual circuit situation, this might lead us to set up a circuit based on maximum capacity paths, but with dynamic routing in the catenet, this is not possible.

Within the current protocols, the best support we can give to aid bulk transfers is to minimise packet overhead using the "Don't Fragment" option. This has the disadvantage, however, that packets must either be dropped, rerouted, or fragmented according to a private protocol if they encounter a net whose packet size is too small (Shoch78). Since the internet protocol makes no guarantee of reliability (Postel78), such packets will most likely be dropped unless the routing algorithm is modified to meet this restriction. That is, for packets with the "Don't Fragment" bit set, the routing algorithm should only consider routes contained entirely within the subset of nets within the catenet which have packet sizes such that fragmentation is not necessary. This clearly has a major effect on the routing algorithm, as it means that the gateways have to keep track of all possible paths, as well as a record of the maximum packet sizes. Moreover, for packets larger than some critical size, there may never be a path in the catenet which is capable of supporting the packets without fragmentation. In this case, the maximum packet size supportable by the catenet on an unfragmented basis is a parameter which must be known by the end processes. In the worst case, in fact, this number may differ depending on the two terminal networks involved.

If the network dependency introduced by the "Don't Fragment" bit ever does become a serious problem, it is most likely to be with bulk traffic, as this service uses the largest possible packet sizes. It is not a general solution to supporting efficient internet transmission. In practise, however, the problem can be mitigated by making use of private fragmentation/reassembly protocols at the right points, which will have the effect of disguising the maximum packet sizes. Where the limiting maximum packet size is governed by a long-haul backbone network, such as the ARPANET, such a private protocol will be absolutely necessary to permit a reasonable end-to-end service.

An alternative is to formalise the existence of gateway-to-gateway fragmentation and reassembly by allocating an option bit "Attempt Reassembly" in the internet protocol, which will be an instruction to gateways to attempt to reassemble the fragments of a packet into larger units, where possible. This option implies that all fragments of a packet must be sent to the same gateway on each internet hop. It also requires that fragments must be stored in the gateway for a certain period before being forwarded, if a gateway decides that reassembly may be possible. Depending on the likely instability of a particular internet route, the first requirement may or may not be a stringent one. The second constraint is more serious in its implications. To be operated effectively it requires that the gateways take steps to avoid congestion and packet loss forcing end-to-end retransmissions, which

would defeat the purpose of minimising packet overheads. In short, it is essential for an "Attempt Reassembly" strategy that there exists effective gateway-to-gateway flow control.

Neither option is entirely satisfactory as a way of supporting efficient transfer in a catenet with the current procedures. The question needs further study.

Conclusions

The three traffic types may be given service support in ways particular to each type. This support may be summarised as follows:

Interactive Traffic	-	Minimise delay
Stream Traffic	-	Minimise congestion
Bulk Traffic	-	Minimise packet overheads

Each of these will directly affect a particular area of the catenet structure, respectively: choice of object function for the routing algorithm; packet queueing techniques inside the gateways; and the fragmentation and reassembly of internet packets. The area in which service support is most unclear is bulk traffic. Further study of the tradeoffs between possible techniques is needed here.

Ultimately, of course, types of service offered will be chosen by the user both on the needs of his application and on his willingness to pay the charges incurred. Because the services required are application dependent, it is important that service types be offered explicitly as minimising delay etc rather than tying them directly to a particular traffic type. In this way the user will be able to choose the combination that best suits his needs. While there is a simple way of providing standard internet priority classes, it is not clear that this can be done so easily with delay classes, while with efficiency it is not clear that anything can be done in terms of grades of service - only an all-or-nothing service has been discussed here.

One point that emerges clearly is that support of internet services depends heavily on the existence of gateway-to-gateway flow control, to produce acceptable rates of packet loss for stream traffic and to permit the tying up of gateway buffers for intermediate reassembly strategies. The need for flow control and exchange of status information between gateways was argued in (Bennett77). iThe high rates of packet loss in both the gateways and the gateway/SIMP VDH links at high data rates in the SATNET gateways (Treadwell78) is directly attributable to the absence of gateway-to-gateway flow control. Such subnet effects can only degrade internet performance still further.

References

Bennett77 - C. J. Bennett, A. J. Hinchley, S. W. Edge, "Issues in the Interconnection of Datagram Networks" IEN 1

Cerf77 - V. G. Cerf, "Gateways and Network Interfaces" IEN 6

Cohen78 - D. Cohen, Private discussion

Postel78 -J. Postel, "Internet Protocol Specification" IEN 41

Shoch78 - J. Shoch, "Inter-Network Fragmentation and the TCP" IEN 20

Strazisar78 - V. Strazisar, "Gateway Dynamic Routing" IEN 25

Treadwell78 - S. W. Treadwell, "SATNET Gateway Calibration Measurements" PSPWN, number to be assigned