

2.3.3.3

Gateways and Network Interfaces

Introduction

The basic issue addressed in this position paper is whether hosts or gateways (if we distinguish between them) should be allowed to selectively refuse traffic from a network.

Historical Perspective

In the early design of the ARPANET, it was recognized that the network might need to summarily block incoming traffic. The bit serial HOST-IMP interface developed by BBN [BBN 1822] had the property that either the IMP or the HOST could block the transmission of bits across the interface. In practice, hosts tried not to block the interface, because IMPs would timeout after a few seconds and discard traffic which they could not deliver to the host. Hosts which were timed out were also declared "dead" to the rest of the network. On the other hand, IMPs were generally free to block incoming (from host) traffic if local conditions (e.g. lack of buffers, multi-packet flow control) required.

An alternative strategy, called the "non-blocking" interface, was pursued in the ARPANET to allow for more flexibility in IMP blocking of incoming traffic. For example, if a host had a multipacket message to send, and the source IMP discovered it needed to reserve buffer space in the destination IMP before the source could send the multipacket message, rather than hard-blocking the HOST-IMP line, the IMP could temporarily refuse the traffic.

Depending on the details of the HOST/IMP protocol, the IMP could request the message later, when it was ready, meanwhile accepting other traffic, or the host could present other traffic for awhile, then present the temporarily rejected traffic again or the host could repeatedly present the rejected traffic until taken by the IMP.

Each of the various strategies for dealing with this kind of interface, have different characteristics and place different requirements on the host and IMP software. For example, if the IMP rejects but wants to ask for the traffic later, the host and IMP must agree on a unique and common name for the message so the host can offer the right one when the IMP asks for it. Alternatively, the host might repeatedly offer traffic from a collection of packets waiting to go, and the IMP would simply reject those it could not serve. This makes the interface between the host/imp slightly simpler, but runs the risk that the host will frequently re-offer traffic which cannot be serviced, or will offer it just after the IMP has timed out the reservation it finally made at the destination IMP to allow the message to be accepted.

Probably, the best strategy is for the IMP and host to agree on a name for the packet and for the IMP to prompt the host when it is able to take the packet. The host can still repeatedly offer the traffic and the IMP can continue to reject, if the host prefers that. The host can discard packets which have been rejected and not requested after a time. Obviously, if the IMP asks for such a message, the host needs to tell the IMP not to bother, or, if the host discards the packet it wanted to send, it could tell the IMP. The IMP could always block the host as a last resort. At issue in this paper is whether hosts should also be able to exercise some form of selective blocking or refusal to accept traffic from a subnet.

Gateways

In the intra-net case, a host exchanging data with another host does not need to rely on the subnet to propagate flow control information to the destination host. The two hosts can protect their respective resources on an end/end basis. The subnet, of course, needs a way to control incoming traffic. For this purpose, it can use blocking or selective blocking, the latter being more flexible.

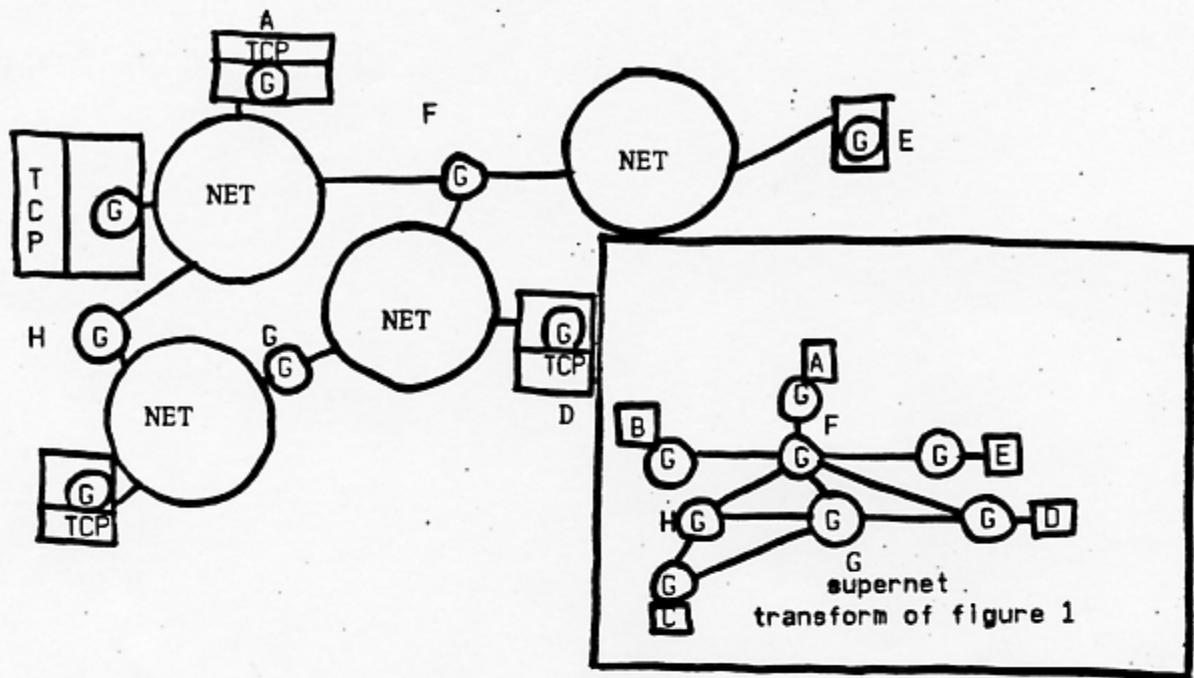
A gateway has always been regarded as a host, at least to the extent that it obeys host/subnet protocols. Gateways which carry traffic out of a subnet to another one do not, in our current design, participate in the host-host (e.g. TCP) flow control mechanism. In fact, the gateway does not presently have any end/end protocol to protect its resources while serving transit traffic.

One possible way to provide for the protection of gateway resources is to imagine a layer of protocol which we might call "gateway-gateway" protocol. Flow control (and perhaps routing) information might pass between any pair of gateways. This model is shown in Figure 1.

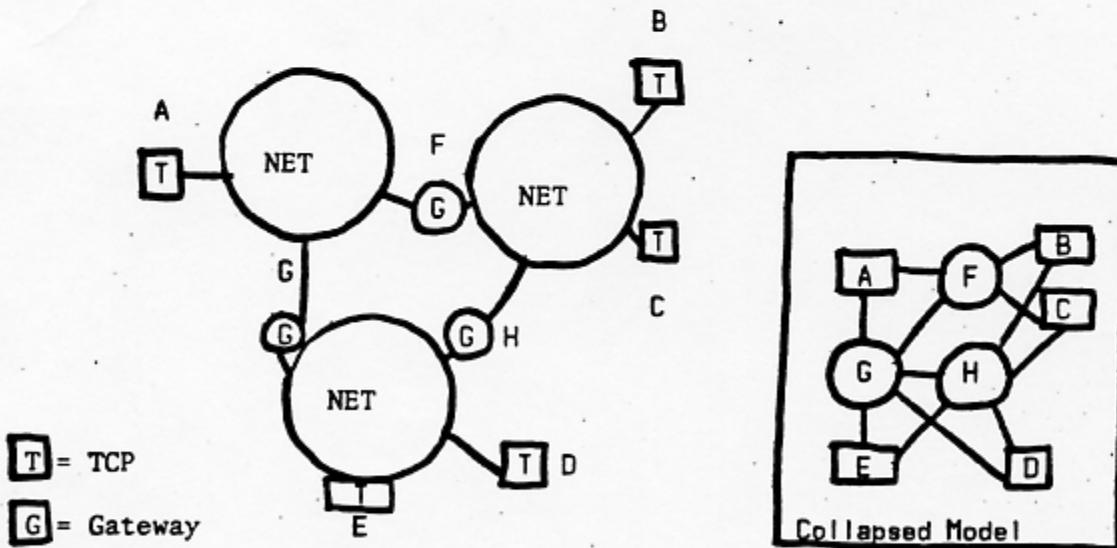
In this model, all adjacent gateways (including those within hosts, associated with TCP's) would exercise flow control with each other.

The disadvantage of this model is that gateways which lie between networks are likely to have a large number of neighbors (thousands!) and to exercise point-to-point flow control could prove costly. Note that this model does not require that the gateways (even those in the hosts) know anything about individual TCP connections. When a particular gateway has a packet to send, it can select a next gateway, independent of the TCP connection involved. Selection of a gateway might be subject to the current gateway-gateway flow control in-force at the moment.

An alternative model is shown in Figure 2.



Gateway Network Model
Figure 1



Transit gateway only model
Figure 2

In this configuration, there is no explicit way for the gateways (labeled F, G, and H) to exercise flow control on TCP's (labeled A-E) through the usual bilateral mechanisms. This is likely to be the case for host resources which cannot supply "gateway" capability. In this instance, it would be highly advantageous for the subnet to accept flow control information from the gateway and to propagate this throughout the subnet so that source packet switches can inhibit the entry of traffic into the net which will not be accepted by the gateway. Having the subnet accept and propagate "selective blocking" information from gateways could substantially simplify the software required in each gateway to deal with internet flow control. Note that this does not rule out further explicit flow control between gateways (i.e. transit gateways joining nets [e.g. F, G, H in Figure 2]).

Furthermore, this facility could be useful for hosts which engage in broadcast protocols, quite independent of the case of a gateway trying to control demand for its resources.

Mixed cases of transit gateways, local gateways (in hosts) and no gateways can also be handled uniformly with the HOST/PS "selective refuse" feature.

Summary

There are other functions which might be found in local host gateways. Generally speaking, we can stick with the conceptual model in Figure 1, but can postulate that some nets will offer sufficient services to hosts that the gateway-gateway flow control can be largely accomplished through HOST/PS selective refusal. In increasing order of convenience or effectiveness, we may obtain Gateway flow control by:

- a. host blocking of channel from IMP
- b. point-to-point flow control between transit gateway and terminating (host) gateway.
- c. crude, distributed flow control through subnet propagation of selective refusal, possibly with explicit flow control between transit gateways.

Discussion

Pure blocking by host or gateway over long periods affects all sources wanting to transmit to that host or gateway. This is not a very effective mechanism for selective flow control. Host or gateway blocking may make sense for purely local and temporary situations where the host or gateway interface to the net is genuinely out of space, quite independent of the existence of a selective refusal service being available from the subnet.

Relying on the existence of "gateways" in each host is dangerous, since some hosts may not have them (or, if they do, they are void of function). The conclusion is that, while gateways should be able to block the HOST/PS Interface (in the PS to host direction), it will be advantageous to consider how gateways (and hosts in general) might report selective blocking or flow control information to an attached net, and how the net might propagate this information to all other attached hosts or gateways. We do not conclude, however, that hosts or gateways should expect the subnets to buffer traffic based on selective refusal by the host to accept a particular incoming packet. This capability is reserved for the host-to-PS direction in which a PS may selectively refuse traffic on a packet-by-packet basis.