

Request for Comments: 852

The ARPANET Short Blocking Feature

RFC 852

Andrew G. Malis
ARPANET Mail: malis@bbn-unix

Bolt Beranek and Newman Inc.
50 Moulton St.
Cambridge, MA 02238

April 1983

This RFC specifies the ARPANET Short Blocking Feature, which will allow ARPANET hosts to optionally shorten the IMP's host blocking timer. This Feature is a replacement of the ARPANET non-blocking host interface, which was never implemented, and will be available to hosts using either the 1822 or 1822L Host Access Protocol. The RFC is also being presented as a solicitation of comments on the Short Blocking Feature, especially from host network software implementers and maintainers.

1 INTRODUCTION

This RFC specifies the ARPANET Short Blocking Feature, which will allow a host to shorten the amount of time that it may be blocked by its IMP after it presents a message to the network (currently, the IMP can block further input from a host for up to fifteen seconds).

The Feature is an addition to the ARPANET 1822 and 1822L Host Access Protocols, and replaces the non-blocking host interface described in section 3.7 of BBN Report 1822 [1], which was never implemented. This Feature will be available to hosts on C/30 IMPs only. This will not present a problem on the ARPANET, which only has C/30 IMPs, but hosts on non-C/30 IMPs in networks that mix C/30 and non-C/30 IMPs will not be able to use the Short Blocking Feature.

The RFC's terminology is consistent with that used in Report 1822, and any new terms will be defined when they are first used. Familiarity with Report 1822 (section 3 in particular) is assumed.

This RFC was once part of RFC 802, which is now obsolete and has been replaced by the combination of this RFC and RFC 851, The ARPANET 1822L Host Access Protocol [2]. The Short Blocking Feature will be available to all hosts on C/30 IMPs, no matter

which (1822 or 1822L) host access protocol they are using to
communicate with the IMP.

2 THE ARPANET SHORT BLOCKING FEATURE

The Short Blocking Feature of the 1822 and 1822L protocols allows a host to present messages to the IMP without causing the IMP to not accept further messages from the host for long amounts of time (up to fifteen seconds). It is a replacement for the non-blocking host interface described in section 3.7 of Report 1822, and that description should be ignored.

2.1 Host Blocking

Usually, when a source host submits a message to an IMP, the IMP immediately processes that message and sends it on its way to its destination host. Sometimes, however, the IMP is not able to process the message immediately. Processing a message requires a significant number of resources, and when the network is heavily loaded, there can sometimes be a long delay before the necessary resources become available. In such cases, the IMP must make a decision as to what to do while it is attempting to gather the resources.

One possibility is for the IMP to stop accepting messages from the source host until it has gathered the resources needed to process the message just submitted. This strategy is known as blocking the host, and is basically the strategy that has been

used in the ARPANET up to the present. When a host submits a message to an IMP, all further transmissions from that host to that IMP are blocked until the message can be processed.

It is important to note, however, that not all messages require the same set of resources in order to be processed by the IMP. The particular set of resources needed will depend on the message type, the message length, and the destination host of the message. Therefore, although it might take a long time to gather the resources needed to process a particular message, it might take only a short time to gather the resources needed to process some other message. This fact exposes a significant disadvantage in the strategy of blocking the host. A host which is blocked may have many other messages to submit which, if only they could be submitted, could be processed immediately. It is "unfair" for the IMP to refuse to accept these messages until it has gathered the resources for some other, unrelated message. Why should messages for which the IMP has plenty of resources be delayed for an arbitrarily long amount of time just because the IMP lacks the resources needed for some other message?

A simple way to alleviate the problem would be to place a limit on the amount of time during which a host can be blocked. This amount of time should be long enough so that, in most circumstances, the IMP will be able to gather the resources

needed to process the message within the given time period. If, however, the resources cannot be gathered in this period of time, the IMP will flush the message, sending a reply to the source host indicating that the message was rejected and specifying the reason that it could not be processed. However, the resource gathering process would continue. The intention is that the host resubmit the message in a short time, when, hopefully, the resource gathering process has concluded successfully. In the meantime, the host can submit other messages, which may be processed sooner. This strategy does not eliminate the phenomenon of host blocking, but only limits the time during which a host is blocked. This shorter time limit will always be less than or equal to two seconds.

Note, however, that there is a disadvantage to having short blocking times. Let us assume that the IMP accepts a message if it has all the resources needed to process it. The ARPANET provides a sequential delivery service, whereby messages with the same priority, source host, and destination host are delivered to the destination host in the same order as they are accepted from the source host. With short blocking times, however, the order in which the IMP accepts messages from the source host need not be the same as the order in which the source host originally submitted the messages. Since the two data streams (one in each

direction) between the host and the IMP are not synchronized, the host may not receive the reply to a rejected message before it submits subsequent messages for the same destination host. If a subsequent message is accepted, the order of acceptance differs from the order of original submission, and the ARPANET will not provide the same type of sequential delivery that it has in the past. If sequential delivery by the subnet is a strict requirement, the Short Blocking Feature should not be used. For messages without this requirement, however, the Short Blocking Feature can be used.

Up to now, type 0 (Regular) messages have only had sub-types available to request the standard blocking timeout, fifteen seconds. The Short Blocking Feature makes available new sub-types that allow the host to request messages to be short blocking, i.e. only cause the host to be blocked for two seconds at most if the message cannot be immediately processed.

Type 0 messages now have the following subtypes:

- 0: Standard: This subtype instructs the IMP to use its full message and error control facilities. The host may be blocked up to fifteen seconds during the message submission.
- 1: Standard, Short Blocking: The IMP attempts to use the same facilities as for subtype 0, but will block the host for a

maximum of two seconds.

- 3: Uncontrolled Packet: The IMP performs no message-control functions, and the packet is not guaranteed to be delivered. The host may be blocked up to fifteen seconds during the packet submission, although any such blockage is unlikely.
- 4: Uncontrolled, Short Blocking: The IMP treats the packet similarly to subtype 3, but will only block the host for a maximum of two seconds. Again, actual blockage is unlikely.

2.2 Reasons for Host Blockage

There are a number of reasons why a message could cause a long blockage in the IMP, which would result in the rejection of a short (or even non-short) blocking message. The IMP signals this rejection of a message by using the Incomplete Transmission (Type 9) message, using the sub-type field to indicate why the message was rejected. The already-existing sub-types for the type 9 message are:

- 0: The destination host did not accept the message quickly enough.
- 1: The message was too long.

- 2: The host took more than fifteen seconds to transmit the message to the IMP. This time is measured from the last bit of the leader through the last bit of the message.
- 3: The message was lost in the network due to IMP or circuit failures.
- 4: The IMP could not accept the entire message within fifteen seconds because of unavailable resources. This sub-type is only used in response to non-short blocking messages. If a short blocking message timed out, it will be responded to with one of sub-types 6-10.
- 5: Source IMP I/O failure occurred during receipt of this message.

The new sub-types that apply to the Short Blocking Feature are:

- 6: Connection set-up delay: Although the IMP presents a simple message-at-a-time interface to the host, it provides an internal connection-oriented (virtual circuit) service, except in the case of uncontrolled packets. Two messages are considered to be on the same connection if they have the same source host (i.e., they are submitted to the same IMP over the same host interface), the same priority, and the same destination host name or address. The subnet maintains

internal connection set-up and tear-down procedures. Connections are set up as needed, and are torn down only after a period of inactivity. Occasionally, network congestion or resource shortage will cause a lengthy delay in connection set-up. During this period, no messages for that connection can be accepted, but other messages can be accepted.

7: End-to-end flow control: For every message that a host submits to an IMP (except uncontrolled packets) the IMP eventually returns a reply to the host indicating the disposition of the message. Between the time that the message is submitted and the time the host receives the reply, the message is said to be outstanding. The ARPANET allows only eight outstanding messages on any given connection. If there are eight outstanding messages on a given connection, and a ninth is submitted, it cannot be accepted. If a message is refused because its connection is blocked due to flow control, messages on other connections can still be accepted.

End-to-end flow control is the most common cause of host blocking in the ARPANET at present.

- 8: Destination IMP buffer space shortage: If the host submits a message of more than 1008 bits (exclusive of the 96-bit leader), buffer space at the destination IMP must be reserved before the message can be accepted. Buffer space at the destination IMP is always reserved on a per-connection basis. If the destination IMP is heavily loaded, there may be a lengthy wait for the buffer space; this is another common cause of blocking in the present ARPANET. Messages are rejected for this reason based on their length and connection; messages of 1008 or fewer bits or messages for other connections may still be acceptable.
- 9: Congestion control: A message may be refused for reasons of congestion control if the path via the intermediate IMPs and lines to the destination IMP is too heavily loaded to handle additional traffic. Messages to other destinations may be acceptable, however.
- 10: Local resource shortage: Occasionally, the source IMP itself is short of buffer space, table entries, or some other resource that it needs to accept a message. Unlike the other reasons for message rejection, this resource shortage will affect all messages equally, except for uncontrolled packets. The message's size or connection is not relevant.

The Short Blocking Feature is available to all hosts on C/30 IMPs, whether they are using the 1822 or 1822L protocol, through the use of Type 0, sub-type 1 and 4 messages. A host using these sub-types should be prepared to correctly handle the Type 9 (Incomplete Transmission) messages from the IMP.

3 REFERENCES

- [1] Specifications for the Interconnection of a Host and an IMP,
BBN Report 1822, December 1981 Revision.
- [2] A. Malis, The ARPANET 1822L Host Access Protocol, Request
for Comments 851, April 1983.

