Hazard-Free Connection Release

Jennifer E. Walter
Department of Computer Science
Texas A&M University
College Station, TX 77843-3112, U.S.A.

Jennifer L. Welch
Department of Computer Science
Texas A&M University
College Station, TX 77843-3112, U.S.A.

Abstract *Fault-tolerant communication in a distributed system requires reliable connection management and message delivery. Reliable connection management includes the guarantee of hazard-free release, in which no data is lost before the connection is terminated. Our work focuses on protocols in which the end nodes retain no connection-specific information between incarnations, operating over networks that deliver packets in order and which may or may not lose packets in transit. In this paper we present a formal model that encompasses the notion of hazard-free connection release. We show that providing a guarantee of hazard-free connection release incurs a penalty over non-hazard-free connection release in terms of message passing overhead if the network does not lose packets. If packet loss may occur, we show that there is no penalty for providing hazard-free connection release, since the connection management protocol must compensate for poorer network behavior.*

Keywords: connection management, distributed computing, graceful release, hazard-free release

1 Introduction

In a distributed computer system, processors communicate over networks which are prone to varying degrees of unreliability, including the possibility that a series of packets will undergo some combination of loss, reordering, and duplication in transit. The goal of a *transport protocol* is to manage connections between hosts, providing reliable communication in spite of a possibly unreliable network layer. The task of managing a connection between two hosts can be divided into: 1) the management of serial incarnations over time, a function known as *incarnation management*, and 2) the control of *message transfer* within an incarnation. Transport protocols are of practical interest because they form the basis for many widely used services such as electronic mail, remote login, file transfer, and remote procedure calls.

Communication protocol specifications must be sensitive to the nature of the network they will operate on and therefore are generally based on some notion of the expected worst-case network behavior and the severity of faults allowed to occur at each processor. Protocols also vary depending on whether the underlying system is synchronous or asynchronous. Synchronous systems can readily incorporate the use of timers for connection establishment and release while asynchronous systems rely on handshaking alone or in combination with connection records and unique connection identifiers to ensure correct message transfer.

In this paper we focus on a particular aspect of connection management – requirements for hazard-free connection release. Hazard-free (a.k.a. safe, graceful, orderly) connection release has been defined informally as a release in which both sides involved in the communication are sure that the data delivery portion of the communication was successful prior to connection release.

This investigation of hazard-free connection release was motivated by the work of several authors [3, 4, 6, 8]. Tanenbaum [6] shows that no handshaking protocol exists to solve the hazard-free connection release problem be-

cause one side will always be in doubt as to whether its acknowledgment was received, and therefore can never disconnect safely. If neither side is prepared to disconnect until it is convinced that the other side is prepared to disconnect too, the release will never take place. However, this proof relies on very poor network behavior, essentially partitioning, in which no packets are successfully delivered after some point. We show in this paper that hazard-free connection release is possible in a (slightly) better behaved network model that guarantees a weak liveness property on packet delivery.

From the results of [1], we know that incarnation management is impossible on non-FIFO, losing networks which return to the same state between incarnations. The results of [2] tell us that incarnation management is impossible over any non-FIFO, losing network in which the nodes have bounded memory. Therefore, in this paper we consider FIFO networks, both losing and non-losing, and prove the precise degree of opening and closing handshake for specific executions on such networks when the hosts have bounded memory and the overall distributed system is asynchronous.

For our results, we assume that no connection records are saved between incarnations, using the definition of *amnesic protocols* from [1]. We assume a connectionless network layer that supports only primitives to send and receive, which does no error checking or flow control, and that produces no duplicate or corrupted packets. However, packets may be lost due to the crash of an intermediate node, packet fragmentation, or congestion on the network. While the FIFO assumption on network delivery may limit the practical application of our results, the lower bounds we prove using this model are theoretically stronger than they would be for weaker network models, providing a first step in the development of a precise characterization of hazard-free connection release. The formal proofs, definitions, and protocols discussed in this paper have been omitted for the sake of brevity and can be found in [7].

An important question we address is how the guarantee of hazard-free connection release influences the cost of a connection, where cost is measured in terms of the number of messages needed to correctly release a connection. This paper represents the beginning of a formal investigation into hazard-free connection release and the ramifications of such guarantees, providing the groundwork for further study.

Section 2 provides formal definitions of the system architecture, incarnation management, and degree of handshake for connection release. Sections 3 and 4 present the handshake requirements for FIFO, non-losing and FIFO, losing networks, respectively. Section 5 provides a discussion of our results and some avenues for future work.

2 Definitions

One of the contributions of this paper is the presentation of a formal model that encompasses the informal notion of hazard-free connection release. This section describes the system modeled and definitions used throughout the paper. The asynchronous system is modeled by I/O automata, introduced in [5]. The remainder of the formalism used to describe the system builds on the model presented in [1].

In this section we define *hazard-free* and *non-hazard-free incarnation management protocols*. Informally, in a hazard-free incarnation management protocol, data may be lost when the connection is terminated and in a non-hazard-free incarnation management protocol, no data is lost when the connection is terminated. We also define *k-way handshake connection release*, which can be informally described as the number of serial packet exchanges between end nodes needed to close the incarnation, measured in a "best case" scenario. These definitions provide the framework for comparing the cost of connection release on network models with various guarantees on quality of service.

3 Handshake Requirements–Non-Losing Networks

This section addresses the degree of handshake required to release a connection on FIFO, non-losing networks in which the protocols are amnesic. Specifically, we show:

1) A one-way handshake release is necessary and sufficient for any non-hazard-free incarnation management protocol.

A straightforward proof shows the necessity of a one-way handshake release, demonstrating that at least one packet must be exchanged upon disconnect or a half-open connection will result. The sufficiency of a one-way handshake release is demonstrated by the presentation and proof of correctness of protocol 1, a non-hazard-free incarnation management protocol. Informally, protocol 1 uses an exchange of two packets to set up the connection, maintaining the connection until one side of the protocol issues a disconnect indication. Since packets are delivered in FIFO order and the network is non-losing, a single packet exchange is shown to be sufficient to correctly release an incarnation.

2) A two-way handshake release is necessary and sufficient for any hazard-free incarnation management protocol.

The necessity of a two-way handshake release is formally proved by showing that packets transmitted at the sender may not be delivered to the receiver prior to disconnect if a one-way handshake release is utilized. Informally, the proof shows that a one-way handshake release can result in an incarnation in which packets are transmitted at the sender but never delivered at the receiver. The proof assumes that a one-way handshake release is sufficient for hazard-free connection release. Let the receiving side of the protocol be called R and the sending side S. An execution is constructed in which the host at R requests disconnect during an open incarnation. Since a one-way handshake is sufficient to close the incarnation, R needs to receive no packets from S after sending a disconnect indication to S and prior to disconnecting. However, this can result in an incorrect execution if the host at S transmits data prior to receiving the disconnect indication from R. Since a one-way handshake is sufficient, R is never alerted that data remains to be sent at S. R may disconnect before receiving this data–a violation of the correctness conditions for a hazard-free incarnation management protocol. Therefore, a two-way handshake release is necessary for hazard-free incarnation management.

Protocol 2, a modification of protocol 1, is presented to show that a two-way handshake release is sufficient for a hazard-free incarnation management protocol. In protocol 2, S is allowed to complete sending pending data prior to disconnect, regardless of which side initiates the disconnect. Protocol 2 therefore meets the correctness requirements for a hazard-free incarnation management protocol.

The main result of this section is the demonstration of a one packet penalty for connection release on FIFO, non-losing networks if hazard-free incarnation management is required. On such a well-behaved network, the incarnation management protocol need not be extremely complex to implement correct connection management. However, added complexity is required to ensure that the connection is hazard-free.

4 Handshake Requirements–Losing Networks

This section examines the handshake requirements for connection release on networks that deliver packets in order, but on which packets may be lost in transit. The results from this section are as follows:

1) A two-way handshake release is necessary and sufficient for any non-hazard-free incarnation management protocol.

The necessity of a two-way handshake release for a non-hazard-free incarnation management

protocol on FIFO, losing networks is proved by contradiction. Informally, the proof goes as follows: Suppose a one-way handshake release is sufficient for a non-hazard-free incarnation management protocol. Then consider an execution where R requests disconnect, but is not required to receive a packet from S prior to disconnecting. Suppose the disconnect indication packet sent from R to S is lost in transit. Since R is not required to receive a packet prior to disconnect, disconnect may occur at R while S has received no notification of the disconnect. It is shown that, in this situation, R must either send packets after it has disconnected to ensure that S eventually disconnects–a violation of the definition of amnesic protocols, or an infinite, half-open connection may result in which S considers a connection to be open after R has disconnected.

The sufficiency of a two-way handshake release is shown by presenting a non-hazard-free protocol, protocol 3, and proving it to be correct. Informally, this protocol works as follows: S and R are synchronized by using the header *first* for packets that indicate the opening of an incarnation. Then the alternating bits 0 and 1 are used to transfer data and synchronize the closing of the incarnation. The sending host requests a connection and S sends a packet with header *first* indicating this intention to R. If R's host is willing to connect, R returns an acknowledgment to S and the connection is open, after which time data transmission can occur. Since the network is losing, all packets are retransmitted by S until they are acknowledged by R. When one of the hosts requests disconnect, a disconnect indication packet is sent repeatedly until it is acknowledged by the other side of the protocol, at which time the connection is complete. The protocol is non-hazard-free because it does not ensure that all pending data is sent after one side requests disconnect and prior to the connection entering a closed state.

2) A two-way handshake release is necessary and sufficient for any hazard-free incarnation management protocol.

The proof of necessity follows from the proof given in section 3 for hazard-free incarnation management on FIFO, non-losing networks. The proof of sufficiency is demonstrated by protocol 4, a modification of protocol 3. Protocol 4 ensures that all transmissions from the sending host have been delivered to the receiving host prior to disconnect, thereby meeting the correctness requirements for a hazard-free incarnation management protocol.

The main result of this section is the demonstration that the one packet penalty for hazard-free incarnation management on FIFO, non-losing networks does not hold for the FIFO, losing case. The poorer quality of service offered by the network in the FIFO, losing case requires the protocol to exhibit the one packet penalty observed in hazard-free incarnation management protocols in section 3 even for non-hazard-free incarnation management. It is somewhat surprising that providing the guarantee of hazard-free incarnation management on FIFO, losing networks requires no additional penalty over FIFO, non-losing networks in terms of message passing.

5 Conclusions

This paper examines the requirements for the degree of opening and closing handshakes in incarnation management protocols which must rely solely on handshake acknowledgment for synchronizing connections, i.e. protocols in which no connection records are retained between incarnations and in which no timing mechanisms are used to measure packet lifetime in the network. FIFO networks are examined with the intention of developing strong lower bounds which are applicable to less well-behaved networks. The results shown are straightforward and rigorously describe the incarnation management process using a system model which is general enough to apply to many existing networks. Thus, this paper provides the groundwork for further research into the requirements for hazard-free connection release under a variety of different network situ-

ations.

We prove that a penalty of an extra packet exchange exists for protocols which ensure hazard-free connection release on FIFO, non-losing networks. A two-way handshake release is shown to be necessary to ensure protocol correctness for both non-hazard-free and hazard-free incarnation management protocols on FIFO, losing networks. Therefore, for FIFO, losing networks, there is no added packet exchange penalty for hazard-free connection release over non-hazard-free connection release. This is because non-hazard-free incarnation management protocols operating on FIFO, losing networks must pay the extra packet penalty incurred by hazard-free incarnation management protocols on FIFO, non-losing networks to ensure correct protocol execution. The possibility of packet loss on a FIFO, losing network requires at least a two-way release to ensure that all packets are delivered in the incarnation in which they are sent and that the incarnations are properly synchronized at both ends of the connection.

The hazard-free incarnation management protocols presented in this paper demonstrate that providing a guarantee of hazard-free connection management is possible on FIFO networks which do not experience crashes of end nodes or partitioning and which do not duplicate packets in transit. Furthermore, providing this guarantee is no more costly on FIFO networks which experience packet loss than on FIFO networks which do not lose packets in transit. Therefore, the problem in providing hazard-free connection release, which was originally raised by Tanenbaum [6], must arise when the network quality of service becomes worse than that offered by the networks examined in this paper, e.g. when the network is non-FIFO and losing.

Future work in this area will involve an investigation of the exact role connection records and timers play in incarnation management over non-FIFO networks with varying degrees of faulty behavior and how the guarantee of hazard-free connection release influences the overall complexity of the communication protocol on these networks. The penalty for guaranteeing hazard-free release on non-FIFO, losing networks will be examined to ascertain the cost incurred by such a guarantee. Since a transport protocol which relies on handshaking alone cannot ensure correct incarnation management on networks which are non-FIFO and losing [1], message passing will not be the sole complexity measure examined on these more poorly behaved networks.

References

[1] H. Attiya, S. Dolev and J. Welch, "Connection Management Without Retaining Information," *Information and Computation*, vol. 123, no. 2, pp. 155–171, 1993.

[2] H. Attiya and R. Rappoport, "The Level of Handshake Required for Establishing a Connection," *The 8th Int'l Workshop on Distributed Algorithms*, pp. 179–193, 1994.

[3] E. W. Biersack and D. Feldmeier, "A Timer-Based Connection Management Protocol with Synchronized Clocks and its Verification," *Comp. Networks and ISDN Systems*, vol. 25, pp. 1303–1319, 1993.

[4] J. G. Fletcher and R. W. Watson, "Mechanisms for a Reliable Timer-Based Protocol," *Computer Networks*, vol. 2, pp. 271–290, 1978.

[5] N. Lynch and M. Tuttle, "An Introduction to Input/Output Automata," *CWI Quarterly*, vol. 2, no. 3, pp. 219–246, 1993.

[6] A. Tanenbaum, *Computer Networks*, Prentice-Hall, 1988.

[7] J. Walter, "Hazard-Free Connection Release," Masters thesis, Texas A&M Univ., College Station, TX, 1997.

[8] R. W. Watson, "The Delta-t Transport Protocol: Features and Experience," *Proc. 14th IEEE Conf. on Local Computer Networks*, pp. 399–407, 1989.