Simulation-based 'STRESS' Testing Case Study: A Multicast Routing Protocol

Ahmed Helmy, Deborah Estrin
Computer Science Department/ISI
University of Southern California
Los Angeles, CA 90089
email:{ahelmy, estrin}@usc.edu *

April 15, 1998

Abstract

Multiparty protocols support an important class of applications ranging from multi-media teleconferencing to network games. Designing wide-area multiparty protocols is becoming more complex with the growth of the Internet and the introduction of new service models. Unexpected combinations of events can drive protocols into undesirable states and may lead to errors. Anticipating all such cases is often impossible and at best may require extensive simulation and testing. In large systems, the cost of testing all possible scenarios exhaustively is prohibitive, and many unexpected cases are not observed until deployment.

Prototyping in testbeds or individual simulations typically focuses on performance under a limited set of randomized protocol transitions. Formal and analytical models representing such protocols tend to be complex, sometimes rendering the model intractable.

In this work, we propose a method for using simulation to analyze the robustness of multiparty (multicast-based) protocols in a systematic fashion. We call our method *Systematic Testing of Robustness by Examination of Selected Scenarios (STRESS).* STRESS aims to cut the time and effort needed to explore the pathological cases of a protocol during its design. This paper has two goals: (1) to describe the method, and (2) to serve as a case study of robustness analysis of multicast routing protocols. We do not prove correctness but aim to offer design method tools similar to those used in CAD and VLSI design. We also demonstrate how powerful and effective systematic simulation can be in studying protocol robustness.

Keywords: *Simulation method, protocol design and testing, robustness evaluation, multicast routing.*

1 Introduction

In this paper, we describe a method for **S***ystematic* **T***esting of* **R***obustness by* **E***xamination of* **S***elected* **S***cenarios (STRESS).* It is based on a simulation framework supported by a set of tools, and is designed for studying protocol behavior in the context of pathological cases and scenarios. Some of the general concepts for STRESS draw from simulation-based verification techniques and reachability analysis [1, 2]. We apply these techniques to support the design, analysis, and testing of multiparty protocols.

In particular, we introduce techniques for state and topology reduction and investigate various packet loss scenarios to capture robustness characteristics. The definition of error conditions enables us to capture the faulty (error-prone) cases automatically.

Multiparty protocols may involve multiple receivers and one or more senders. These protocols include multicast routing protocols (e.g. DVMRP [3], MOSPF [4], PIM-DM [5], CBT [6], and PIM-SM [7]),

*This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. DABT63-96-C-0054. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the DARPA.

Figure 1: Establishing multicast delivery tree

Our study revealed several pathological errors in PIM-SM through the use of STRESS tools, and evaluated solutions to eliminate these errors. The suggested solutions have since been added to the PIM-SM specification [15].

The rest of the paper is organized as follows. Section 2 provides an overview of the STRESS method. The case study for PIM-SM is presented in section 3. Results are given in section 4. Sections 5 and 6 address related work, summary and future work, respectively.

2 The Method Overview

The *robustness* of a protocol is its ability to respond correctly in the face of network failures and packet loss. The goal of the *STRESS* method is to provide a framework for systematic testing of protocol robustness through the examination of selected scenarios.

For a given protocol, we first capture a set of error-prone scenarios. This is achieved by: (a) investigating a *representative* subset of the protocol state space, and (b) defining error conditions. We use these scenarios to evaluate design trade-offs, analyze behavior, and as a test suite to examine various implementations of the protocol.

Our basic approach consists of three stages: *scenario generation* (pre-processing), *tracing* (simulation), and *output analysis* (post-processing). Figure 2 illustrates these stages. The building blocks in the figure are explained in detail throughout the rest of this section.

Note that the engineering design process is usually iterative, where an investigator may cycle and feedback into previous stages based on his/her intuition and insight, sometimes gained by the analysis of earlier simulations. Our methodology does not contradict such process. In fact we will show in the results section how we iterated through the stages to guide our simulations. This section, however, only discusses the modules supporting the different stages.

2.1 Scenario Generation

Scenarios are composed of routed topologies and sequences of events (input stimuli and state transitions), and describe the simulation context that may cause protocol transitions. Scenario parameters include the *routed topology*, *host scenarios* and *loss scenarios*.

2.1.1 Routed topology

The routed topology is the network infrastructure upon which the protocol operates; e.g. network nodes and links and the unicast routing that determines how packets are forwarded.

We try to identify simple topologies that capture a large percentage of the protocol's state space, and to which other, more complex, topologies may be reduced[1]. For much of our study, we choose a LAN with four connected routers as the basic topology. We show how other topologies are reducible to the *four-router* LAN topology, and discuss the limitations of such a topology in section 3.2. We further extend the topology to capture particular characteristics of the protocol under study, PIM-SM.

As a component of the routed topology, unicast route inconsistencies may be a common source of error. Unicast routing may exist in one of the following three states: (a) consistent routing, (b) transient inconsistent routing, and (c) long lived inconsistency. Case (a) requires no changes. The study of case (b) is convergence analysis, which has been addressed

[1] Two topologies are said to be reducible (or equivalent) if they drive the protocol (according to the host scenarios applied) into the same states, experiencing the same set of state transitions.

Figure 3: The scenario filter

For all possible permutations, there exists $5! = 120$ scenarios, considering that each host event occurs once. Then, as shown by figure 3, we apply protocol constraints, (e.g. *a receiver cannot leave before it joins the group*), to reduce the number of possible combinations to $5!/(2! \times 2!) = 30$ scenarios. Further, as a practical input, we assume, without loss of generality, that (*the source sends packets throughout the simulation*), to reduce the number of possible scenarios to $30/5 = 6$ scenarios. These six scenarios are:

`1. J1:J2:L1:L2 2. J1:J2:L2:L1 3. J1:L1:J2:L2`

`4. J2:J1:L1:L2 5. J2:J1:L2:L1 6. J2:L2:J1:L1`

The number of representative scenarios can be even reduced further if the host distribution is sym-

metric with respect to the topology, since the following scenarios will be equivalent: (i) 1 equivalent to 5, (ii) 2 equivalent to 4, and (iii) 3 equivalent to 6; i.e. we need only investigate 3 different host scenarios for the given topology.

These scenarios may be generated automatically by the method. However, generalizing the process of obtaining representative scenarios for various multiparty protocols is currently under study[4].

2.1.3 Loss and Failures

The loss and failure scenarios considered include the loss and corruption of packets (during transport, routing or forwarding within the network), or loss of state in router nodes due to protocol daemon failures, machine crashes, or insufficient resources (e.g. memory).

In this study, we address single fault models; i.e. those that address the occurrence of a single fault per scenario (e.g. single packet loss).

Loss of packets Packet loss may occur in various parts of the network, due to congestion, or link, node or interface failures. We classify these events as simply 'packet loss' regardless of cause, and create exhaustive loss scenarios to capture all the possible protocol transitions and pathologies due to packet loss.

Here, the fault model used for packet loss is the loss of a single protocol message between the sender of that message and any of the intended receivers.

For most multicast protocols, when routers are connected via a multi-access network (or LAN)[5], hop-by-hop messages are multicast on the LAN, and may experience selective loss; i.e. may be received by some nodes but not others. The likelihood of selective loss is increased by the fact that LANs often contain hubs, bridges, switches, and other network devices. Selective loss may affect protocol robustness. Similarly, multiparty protocols and applications must deal with situations of selective loss. This differentiates these applications most clearly from their unicast counterparts, and raises interesting robustness questions.

Our case study illustrates why selective loss should be considered when evaluating protocol robustness. This lesson is likely to extend to the design of higher layer protocols that operate on top of multicast and can have similar selective loss.

The input to the 'loss & failures' substage (shown in figure 2) is obtained from initial traces of simulations without protocol message loss. These traces guide further simulations to cover all possible protocol message loss scenarios.

Loss of state State loss in nodes may occur due to crashes, loss of multicast, unicast or all forwarding entries. We investigate how such loss affects the protocol's overall correctness, especially from the end-systems perspective.

2.2 Simulation and Tracing

During this stage the protocol mechanisms are simulated and traces are collected:

2.2.1 Simulation

One desirable approach for simulating complex protocols, is to include detailed mechanisms of parts of the protocol while abstracting out others, we call this approach *subsetting*.

Subsetting refers to selecting subsets of the protocol functions, while abstracting or removing others. This allows us to focus on specific parts of the protocol state space. Subsetting can be based on:

- *Protocol functions*: Subsetting protocol functions (or mechanisms) refers to the abstraction of these functions. This may be achieved by replacing a complex mechanism by a simpler one, exhibiting similar external behavior under relaxed assumptions. For example, one may use static configuration instead of simulating a detailed bootstrap algorithm. This way, one may study other protocol mechanisms assuming correctness of the bootstrap mechanism.
- *Protocol states*: A study may focus on specific protocol states. This allows, for example, the study of multicast group state without dealing with source-specific state.
- *Messages types*: This allows the examination of specific protocol message types in the absence of others.

Note that subsetting does not permit all combinations of the above items. To maintain protocol correctness, an abstracted part has to be replaced by its equivalent that exhibits similar external behavior. For example, one may not simply remove the bootstrap mechanism and run the simulations.

[4] This document does not attempt to address automatic scenario generation, which is the subject of our on-going work. We are investigating two approaches –based on forward and backward search techniques– to extract error-prone scenarios and topologies automatically, from a finite-state-machine model of the protocol.

[5] We use the term LAN to designate a connected network with respect to IP-multicast. This includes shared media (such as Ethernet, or FDDI), hubs, switches, etc.

2.2.2 Tracing

Tracing is the process of logging information about events or packets during the simulation run. Logged information is analyzed during the post-processing (i.e. the output analysis) stage. In addition, some traces are used as feedback to the scenario generator to guide further simulations. We consider several kinds of tracing:

End-point tracing Tracing end-points includes logging information pertaining to hosts sending or receiving packets, and joining or leaving multicast groups. A detailed description of the traces used in the case study is given in section 3.

To identify errors and pathologies in the protocol itself, we focus on the effect of the multicast routing protocol transitions on the end-point packet delivery (as explained in section 2.3).

Protocol state transition tracing A protocol can be represented by a finite state machine (automaton), consisting of states, transitions and stimuli (inputs, outputs, and timer actions). Based on knowledge of initial protocol states, we obtain the sequence of protocol transitions by tracing all stimuli.

We use protocol traces to diagnose and verify protocol behavior, and to analyze errors.

Link tracing We keep track of packets traversing links (or LANs) between nodes, as well as the events of packet loss on links (or LANs). This information is used in several ways: in output visualization, output analysis, or as feedback for scenario generation. Links carrying message types of interest are targeted for intentional packet loss in further simulations. This reduces the number of loss scenarios examined to those directly affecting the protocol behavior under investigation.

Code annotation When placed in key points (such as beginning of protocol procedures or code modifying the state of the protocol), code annotations capture internal execution of the protocol machinery. We use code annotation to estimate what part of the code, and subsequently the protocol, has been executed and stressed (code coverage).

2.3 Output Analysis

One major concern of STRESS is to identify pathological cases, and indicate when and if an error occurred and why. This is achieved in the output analysis stage, which consists of:

Identifying end-point errors Error conditions may be specified with respect to end-point traces. Examples of end-point error conditions are black holes[6], and packet duplication, where more than one copy of the same packet is received by a group member.

If the factors during one simulation run are relatively static (i.e. static unicast routing, static topology and controlled loss), the error may be attributed to an error in the multicast routing protocol.

Once the specified error is identified by the output analyzer, the trace log is rolled back in time to investigate the protocol traces, as explained next.

Relating errors to protocol After detecting an end-point error, the output analyzer isolates the possible causes of such errors in the form of protocol traces. The output analyzer in this case is similar to a logic analyzer, allowing the designer to navigate backward in time and investigate the causes of the error.

As will be shown (in section 3), the process of identifying a protocol error may suggest fixes to the problem.

Code profiling The profiler captures information about the annotated code, such as which procedures were (or were not) invoked, and the order and frequency of invoking protocol procedures. This information indicates the portion of the protocol stressed by the examined scenarios.

3 Case Study of PIM-SM

To evaluate the utility of STRESS, we applied it to a complex multicast routing protocol; PIM-SM. Before going into details of the case study, we first give an overview of multicast routing and PIM-SM.

3.1 Multicast Routing – Overview

Multicast distribution trees may be established by either broadcast-and-prune or explicit join protocols. In the former, such as DVMRP or PIM-DM, a multicast packet is broadcast to all leaf subnetworks. Subnetworks with no local members for the group send *prune* messages towards the source(s) of the packets to stop further broadcasts. Link state protocols, such as MOSPF, broadcast membership information to all nodes. In contrast, in explicit join protocols, such as CBT or PIM-SM, routers send hop-by-hop *join* messages for the groups and sources for which they have local members. When received, these messages

[6] A black hole is an observable amount of consecutive packet loss between periods of packet delivery.

Figure 5: The equivalent topologies

3.2.1 Prunes

First, we consider *N-router* LAN topologies, where $N = 1, 2$, and 3, respectively. It is trivial to prove

[7] To achieve this, a downstream router receiving a *prune* on its incoming interface triggers a *join* upstream.

that these topologies are not equivalent for hop-by-hop messages.

Assumption *N-router LAN topology, where $N > 3$, is reducible to the three-router LAN topology for prunes, w.r.t. single message loss scenarios.*

To justify our assumption we first prove that a *four-router* LAN topology is reducible to a *three-router* LAN topology.

Correctness condition: *If a router on the LAN has the LAN as its incoming interface, there must be one other router with the LAN in its outgoing list.* Once this condition is satisfied, violating it is considered a protocol error[8].

Next, we examine the *three-router* LAN topology. In figure 5, topology 'I', assume that A and B are downstream routers, and C is an upstream router.

- In figure 5, topology 'I', router C has the LAN in its outgoing list, router A has the LAN as its incoming interface, and router B is leaving the group and so sends a *prune* towards C. The *prune* is multicast on the LAN.

  The only case where the correctness condition may be violated is when C receives the *prune* while A does not. In the other cases, either the *prune* is not received by C, or is received by A which triggers a *prune-override* to re-establish the LAN in C's outgoing list. This is illustrated by the following selective loss pattern table for the *prune* message sent by B:

A	C	
0	0	
0	1	
1	0	← error
1	1	

  where a '0' indicates no-loss and '1' indicates loss. The error occurs where the upstream router (C) received the *prune*, but the router with downstream members (A) did not receive it.

- In figure 5, topology 'II', we add another downstream router D. The selective loss pattern table follows:

A	D	C	
0	0	0	
0	0	1	
0	1	0	
0	1	1	
1	0	0	
1	0	1	
1	1	0	← error
1	1	1	

  The only error occurs when the upstream router (C) receives the *prune*, but neither of the downstream routers receives it. If the *prune* is received by any of the downstream routers, a *prune-override* would re-establish the LAN in C's outgoing list.

From the symmetry of the loss patterns and topology we see that all errors are triggered by the same transitions experienced by router A in topology 'I'. Hence, the extended topology 'II' does not introduce any new errors, and exhibits the same external behavior as does topology 'I'. We conclude that topology 'I' and topology 'II' are equivalent for *prunes*.

We now show that the *N+1-router* LAN topology is reducible to the N case; where $N \geq 3$.

With the addition of an upstream router (figure 5, topology 'III'), no added error cases are encountered. The addition of a downstream router, however, may introduce new error scenarios. Similar to the *four-router* LAN case, we establish the following assertion: *'the only error case occurs when all downstream routers lose the prune and the upstream router receives it.'* If the *prune* was received by any of the downstream routers, the correctness condition would be retained using *prune-overrides*[9].

The assertion holds in both topologies. Hence, we conclude that the *N-router* topology experiences the same errors as the *N+1-router* topology.

From the above we see that by simulating the *three-router* LAN topology we capture all the errors, with respect to selective loss (for the *prune* mechanism), that may be experienced by any *N-router* LAN topology; where $N > 3$.

3.2.2 Joins

For loss-free scenarios, the equivalence proof is straightforward. For lossy scenarios, the loss of a *join* mes-

[8] This is to differentiate between join latency (which is not considered a protocol error) and a black hole which is a protocol error.

[9] prune-overrides are actually join messages. The effect of join message loss is described in section 3.2.2.

Figure 6: The topology used for the case study

For our case study, we use a *four-router* LAN topology with an added Rendezvous-Point (RP) to capture shared tree characteristics. The overall physical topology consists of five routers, four of which are connected via a LAN, as shown in figure 6.

3.3 Test suites

In this section we elaborate on the routed topology, host scenarios and loss pattern generation used for our case study. We also describe the simplifications and subsettings applied.

[10] Note that if join suppression as described by the PIM-SM specification [15] is implemented, the equivalence relation can be established similar to the *prunes* case above.

[11] A limitation to the *four-router* LAN topology is given for the esoteric case of three upstream routers and three downstream routers with inconsistent unicast routing tables. This case creates one extra transition that can only be captured by at least a *six-router* LAN topology. We do not consider this a practically significant scenario, and we consider its analysis as a special case, not captured by the *four-router* LAN topology.

However, aside from this exception, the *N-router* LAN topology (where $N > 4$) is equivalent to a *four-router* LAN topology w.r.t. *asserts*.

Physical and routed topologies The overall topology used is that shown in figure 6. For the unicast routing protocol we use a centralized version of Dijkstra's Shortest Path First (SPF) algorithm [17].

PIM-SM uses the underlying unicast routing tables for building multicast trees. Therefore, unicast routing inconsistencies affect the operation of PIM-SM. To investigate such interaction we add a component to force inconsistent multicast routes between PIM routers, as shown in figure 6, topology 1. Inconsistent unicast routing cases arise, for example, when a multicast region spans multiple unicast domains (or ASs)[12].

Host scenarios Since protocol states for different groups do not interact, we consider only one group. Also, since protocol states for different sources do not interact, it suffices to consider only one source 'S' per simulation run[13]. The source is modeled as a constant-bit rate (CBR) stream with fixed packet size. The source model does not affect the correctness of the method. However, to assure full controllability over the selective loss model, we set the data rate to ensure that no loss occurs due to queue overflow[14].

While we consider only a single source, we consider two receivers ('R1' and 'R2') for the same group to account for shared tree state interactions. We use the host scenarios described in section 2.1.2.

Loss patterns We investigate all possible selective loss scenarios for multicast hop-by-hop PIM-SM messages in this representative topology.

Loss models are applied exhaustively to those links that carry the protocol messages under investigation. The tracing stage identifies these links during the first simulation run, without packet loss, and feeds back the link information to the loss generation module, as shown in figure 2. As we will show in section 3.5, the number of representative scenarios is quite small, and hence the number of overall lossy scenarios explored is manageable.

We do not address state loss or node crashes in this document. However, crash scenarios may be implemented in a way similar to loss scenarios.

[12] A similar interaction with unicast routing may be created easily when switching to the shortest paths in PIM-SM.

[13] We do not consider aggregated source or group entries in this study.

[14] For this we use packet size of 180 bytes, and a send interval of 25 ms (i.e. source rate of 57.6 kb/s), this ensures no queue drops on the 1.5 Mb/s links used with 10 packet queue limit.

Tracing Trace information includes the event type (send or receive), the node experiencing the event, the type of message sent or received, and the time at which the event occurred. Every data packet is assigned a unique sequence number.

For example, a trace may take the following format: '**R2 Node A Rcv 7 t 190**' meaning that receiver **R2** in node **A** received a data packet with sequence number **7** at time **190ms** from the beginning of the simulation run.

Subsetting For brevity, we do not consider source-specific trees and switching to the shortest paths in this paper. This is an example of *state subsetting*, since we consider shared group states while disregarding source-specific states.

The messages considered in the study are *join*, *prune*, *assert* and *register* messages. To study *joins*, *prunes* and *asserts* without the effect of *registers*, we consider a topology where the source and the RP are co-located (see S1 in figure 6, topology 1). This is an example of *message subsetting*.

When studying *registers*, *joins* and *prunes* we consider topology 2 in figure 6 where: (a) S2 is the source, hence node A sends registers to the RP, and (b) the routed topology has consistent unicast routing, to eliminate the effect of the *assert* mechanism. This represents *function (or mechanism) subsetting*.

Only triggered actions are investigated for simplicity. Timer action analysis is not considered in this study, and is part of our work-in-progress.

3.4 Applying the Method

This section provides an illustrative example, showing how STRESS may be used to identify and analyze errors encountered during the simulation of the representative scenarios.

3.4.1 The Simulation Framework

We have implemented an initial version of the STRESS method in the Network Simulator 'NS' [18]. NS is an event-driven packet-level simulator controlled and configured via Tcl [19] and Object-Tcl (or OTcl [20])[15]. To support our method, we have added modules to provide LAN support, controlled selective loss, protocol tracing, profiling capabilities, and a detailed implementation of PIM-SM[16]. This implementation serves as the simulation environment for our case study. In addition, the building blocks were designed to be re-used within the same framework to apply this method to other multiparty protocols.

Figure 7 gives a high level description of the network and protocol simulation modules, described below.

Network simulation The network model consists essentially of links and nodes. The building blocks (e.g. the queues and the forwarding machineries) are implemented in *C++* as part of NS core mechanisms, to improve per packet processing performance. These building blocks are connected using OTcl objects, such as links, LANs and network nodes.

Link A link consists of the following building blocks:

- *Queue* structure provides enqueuing, dequeuing and drop primitives. The queue drop policy is selected from various options (e.g. drop-tail, or RED [22]).
- *Delay* module simulates the transmission and propagation delays on a link.
- *Loss* module controls the loss patterns on the link. The loss may be based on the type of message according to a loss-trace file.
- *Tracer* modules tap various components of the link, to capture packet and loss trace information. These modules provide full observability of the link events.

LAN A LAN is built of multiple links in addition to a *replicator*. When the replicator receives a packet, it copies and delivers it to each of the connected queues and subsequently, to neighboring nodes. This model facilitates the simulation of selective loss, as the loss modules for the connected nodes are controlled independently.

Network node The network node provides the multicast and unicast forwarding mechanisms, independent of routing protocols. This model resembles

[15] For information about the simulator see http://catarina.usc.edu/vint.

[16] Our detailed PIM-SM simulation mimics the unix 'pimd' [21] implementation model, and hence is able to capture many implementation aspects. We plan to develop an interface between the simulator and an operational network running the pimd code. However, the analyses presented in this study are based strictly on the protocol specification, independent of the implementation.

Figure 7: Network and Protocol simulation building blocks

that of a unix kernel forwarding mechanism. Following are the basic building blocks for the network node:

- *Packet classifier* receives packets from the network interfaces or the local protocols and passes them to the unicast or multicast forwarding mechanism as appropriate.
- *Multicast and unicast forwarding* provide the lookup tables used to forward multicast and unicast packets. The forwarding tables at a node are installed by the routing protocols running at that node.

Protocol simulation The protocol was implemented using the dynamic object-oriented language OTcl, which provides flexibility and rapid prototyping. Protocol modules were designed to facilitate subsetting of functions.

Architectural modules A protocol instance is attached to an NS network node, and performs the protocol functions specified. The protocol consists of the following architectural building blocks:

- *Send and receive primitives* –also called *message agents*– handle incoming messages dispatching to the appropriate processing module, and sending messages to the unicast or multicast forwarding machinery within a node, or directly to the network interface. One may think of these primitives as a socket layer.
- *Send handlers* encode the outgoing messages into the proper packet format and pass the messages to the send primitives. Each send handler handles a specific protocol message type (for example, Prune).
- *Receive handlers* process the incoming protocol control messages. Each handler processes a specific protocol message type.
- *Protocol state storage* represents the state of the protocol in terms of entries, flags and timer values. An example of state storage is the multicast routing tables.
- *Timer handlers* process the protocol timed actions upon expiration of timer values kept in the protocol state.

3.4.2 Simulating PIM-SM

To verify that our PIM-SM implementation conforms to the protocol specification, we ran several conformance test-suites using the simulator.

Obtaining faulty scenarios To obtain the *faulty scenarios* (i.e. those that contain errors), we execute the method stages in order (i.e. scenario generation, simulation and tracing, and output analysis, respectively), and then reverse the order from the output

Figure 8: Simple packet trace graph showing packet loss and duplication

to the traces to identify the faulty scenarios. These phases are automated by the tools provided, and are transparent to the user, once the scenario set-up is complete.

The process of attributing end-point errors to protocol actions may be automated only if the error conditions are given in terms of such protocol actions. In practice, however, these protocol error conditions are often not known *a priori* by the designer, and are usually defined in terms of end-point errors (such as packet loss or duplication). The supporting tools identify end-point errors, and provide a history of protocol traces. The designer then examines the traces and identifies the protocol errors[17]. This process may suggest fixes to the problem, as we will show in the results section.

Example In our simple example, an error condition is any packet loss or duplication, experienced by the end-points. A faulty scenario (without packet loss) that leads to two error conditions, is identified and explained. Then the protocol actions leading to the errors are analyzed.

The representative scenario explained here is 'J1:J2:L1:L2' using topology 1. This scenario was identified automatically as a faulty scenario. Traces in figure 8 give the history of the errors found. The first error (i.e. the packet duplication) has the host event 'J2' as the closest join or leave host event in its history at time 200ms. The error is a join transient caused by parallel paths to the RP. The error is resolved using the *Assert* messages exchanged during the duplication at time 246ms. The second error (i.e. packet loss) is a leave transient; it has a host event 'L1' in its recent history. The loss is due to the *prune* sent by node A at 300ms, and is resolved by a *prune-override* sent by node B at 310ms.

Although the protocol actions leading to the end-point errors (specified as any packet loss or duplication in this specific example) are considered transient errors, they are not considered protocol design errors. We do, however, address protocol design errors in section 4.

3.5 Scenario and protocol coverage

While the fact that we were able to discover design errors provides some evidence of the method's utility, we would like to quantify the coverage of protocol states and possible scenarios.

The overall protocol coverage has two dimensions. The first is the protocol state coverage, and we attempt to cover this dimension using the *representative* scenarios' reachable states. Investigation of the *loss* scenarios does not affect protocol coverage significantly.

The second dimension is the space of possible interaction scenarios between these state machines (in different routers) within the topology. This dimension is explored by investigating the *selective loss* scenarios.

[17] In our limited experience, we have identified protocol errors in the recent history of the end-point errors.

Scenarios covered The initial number of simulated scenarios *without* protocol message loss was:

$$\sum_{topologies} \text{(No. rep. scenarios)}$$

Where No. rep. scenarios is the number of *representative* scenarios, equal to 6 in our case (discussed in section 2.1.2), and the topologies are the two discussed in section 3.3. Hence we simulated 12 scenarios without protocol message loss.

After feeding back the link traces for the messages under study, the loss patterns were assigned to the corresponding links. The scenario generator then setup the simulations for the new scenarios with loss.

The total number of scenarios *with* protocol message loss simulated is given by the following formula:

$$\sum_{\forall Topos} \left(\sum_{\forall Reps} \left(\sum_{\forall Msgs} \left(\sum_{\forall Links} LinkMsgs \cdot 2^{(LinkRtrs-1)} \right) \right) \right)$$

where the terms used are described in the following table.

Term	Meaning
Topos	Topologies
Reps	Representative Scenarios
Msgs	Messages under study
LinkMsgs	No. messages traversing the link
LinkRtrs	No. routers connected to the link

For each topology, this formula gives the number of scenarios automatically generated after the first simulation run, during which the number of messages and links (traversed by these messages) is counted.

For example, for the first topology, the messages under study were *joins*, *prunes* and *asserts*. The representative scenarios triggered 16 *joins*, 12 *prunes*, and 12 *asserts* on the LAN, and 16 *joins* and 16 *prunes* on point-to-point links. Hence, the total number of scenarios with loss became:

$$(16 + 12 + 12) \cdot 2^3 + (16 + 16) = 352 \text{ scenarios},$$

For the second topology, the messages under study were *joins* and *prunes*. The representative scenarios triggered 16 *joins*, and 18 *prunes* on the LAN, and 6 *joins* and 18 *prunes* on point-to-point links. Hence, the total number of scenarios with loss became:

$$(16 + 18) \cdot 2^3 + (6 + 18) = 296 \text{ scenarios},$$

Protocol code coverage A large portion of the multicast support code in NS was annotated automatically to provide code tracing. Out of 91 procedures (procs), following is a summary of the procedures covered by the representative scenarios:

Topology	Procs covered	%ge
Topology1	79	86.8%
Topology2	80	87.9%
Total	84	92.3%

Procedures that were not invoked dealt mainly with source-specific state (which was abstracted in our test suites), or with the modularity of the object-oriented nature of the code.

4 Results

This section describes the protocol design errors uncovered for PIM-SM under STRESS.

Unlike our simple example above, we are only interested in design (i.e. not transient) errors. For this, we modified the error conditions to avoid join and leave transients. The new error conditions do not consider single duplication or loss.

Following is a summary of the major faulty scenarios encountered, and how they relate to STRESS. For a more detailed discussion of the protocol errors and fixes see section 4.2.

4.1 Summary of Results

We describe a partial list of *faulty scenarios* captured by STRESS. We obtained this list after simulating only a few of the representative scenarios. The traces produced provided guidance to discover the protocol errors. Design errors discovered include *Assert*, *Join/Prune* and *Register* mechanisms.

Asserts For the first topology (figure 6, topology 1), a black hole was observed for one receiver.

The faulty scenario in this case involved another receiver joining in the recent history of the black hole. By analyzing the protocol trace history after rolling back, we noticed that an *Assert* process took place right before the loss.

In addition, the faulty scenario included the loss of a *join* message, which prevented the establishment of the branch of the shared tree from the Assert winner to the RP. Hence, the protocol design error is allowing a router on a branch of the tree that is not completely established, to participate in *Asserts*.

Joins and Prunes Over the same topology (i.e. figure 6, topology 1), several other faulty scenarios lead to black holes. The host scenarios involved one receiver leaving just before black holes were experienced by the other receiver. In these cases *join* and *prune* messages occurred the recent history of the end-point error.

Furthermore, all such scenarios included either: (i) loss of a *join* message, preventing a pruned branch from being re-established; or (ii) selective loss of a *prune* message, preventing a *join* (i.e. *prune-override*) from being triggered. The protocol design error in this case was not allowing a second chance for routers with downstream members to override *prunes*.

Registers In the second topology (figure 6, topology 2), faulty scenarios were captured that cause packet duplicates at the end-points.

In this case, the observed faulty scenarios did not follow a regular pattern, and were developed iteratively (i.e. when one faulty scenario led to a suggested fix in the protocol, the fix was implemented and the method re-run to observe further faulty scenarios).

The first scenario involved a single host receiving duplicates merely by joining the group. The packets were being delivered at least twice, once directly from the source –by virtue of being on the same LAN–, and the second delivery from the shared tree after the *register* reached the RP and was sent down the shared tree. When the number of packet duplicates exceeded two, this suggested a loop. The loop occurred when a packet received over the shared tree on the LAN, was (a) picked up by the local router, (b) re-registered to the RP, and (c) forwarded down the shared tree again. The protocol error was allowing the packets to flow down from the shared tree to the originating LAN, and be re-registered. The fix was to prune such sources from the shared tree.

The second scenario involved another receiver joining before the duplicates were observed. The pruned branch of the shared tree was re-established by the joining receiver, allowing the packets to flow down the shared tree to the originating LAN, and subsequently, causing the loop.

The third scenario involved a *prune* message loss, again allowing the packets to flow down the shared tree to the originating LAN, and led to looping.

Rules were added to prevent packets from being forwarded back on their original LANs in the above scenarios.

4.2 Detailed Results

The rest of this section describes the above faulty scenarios in more detail, and illustrates how the solutions were developed with the aid of STRESS. After the solutions were integrated into the protocol simulator, we applied 'regression testing' to verify that the fixes did not introduce any new errors.

4.2.1 Assert analysis

Following is a discussion of the pathological cases found in the *Assert process*. An exhaustive list of the results is not included in this document for brevity. A few errors in the PIM-SM specification were unveiled during this process, we focus on errors that created the possibility of packet loss (i.e. black holes).

Figure 10: Transition diagram for *join*s and *assert*s

4.2.2 Join/Prune analysis

In this analysis we address the effect of selective loss of *Join/Prune* messages. Although this problem has been addressed in recent releases of the PIM-SM spec-

ification, we provide a more efficient solution.

We use the topology given in figure 11, **I**. The representative scenario used is 'J1:J2:L1:L2' with the second *join* from node A lost on the LAN.

We assume that S1 sends packets to group G throughout the simulation. Consider the sequence of events given in figure 11, **I**. After the last event (step 5), R2 stops receiving S1's packets. This problem persists until A sends the next periodic *join* to C and re-establishes the pruned branch of the tree. A similar problem is encountered in figure 11, **II**, when the *prune* sent from B is selectively lost on the LAN by A and received by C.

Discussion and fix The solution suggested by the PIM specification introduces a deletion timer. This, however, increases the leave latency, and incurs unnecessary data overhead.

A more efficient solution would be to have the upstream router (C) announce a *'prune-alert'*, before removing the LAN from its outgoing list, by resending the *prune* message previously received from B.

4.2.3 Register analysis

Following is a description of the scenarios that exhibit packet duplication due to *register* messages, and the suggested fixes to eliminate such duplication. The fixes were applied iteratively, until the error was eliminated.

First scenario (single source, single receiver): In this scenario we consider S2 and R2 in figure 12 **(I)**. Consider the sequence of events in the figure.

Packet duplication and *register* looping occur in the above scenario. A similar scenario occurs when R2 joins first then S2 starts sending to the group.

Suggested fixes The required behavior is to send a –triggered and periodic– source-specific *prune* off of the shared tree, if a router has source-specific state for registering and shared tree state for the same group (regardless of the incoming interface settings).

Second scenario (single sender, two receivers): We assume the implementation of the above fixes to the simulator, then consider the sequence of events in figure 12 **(II)**. This scenario exhibits packet duplication and *register* looping.

Suggested fix The problem arises because the packets are forwarded back on the originating LAN, and treated as if they were new packets originated by the directly connected source. The following rule solves this problem for the given scenario:

- A router receiving *join* message must **NOT** add an interface on the same subnet as a source S, for any source specific entry for S associated with same group.

Third scenario (single source, single receiver with message loss): Considering the scenario in figure 12 **(III)**.

The source specific *prune* sent from A to C –when A having a shared tree state, creates the source specific entry for registering– is lost.

Packet duplication and *register* looping problems are experienced in this scenario. The problem persists until a periodic *Join/Prune* message is successfully sent upstream.

Suggested fix To be robust to (at least) one message loss, we suggest the following rule for packet forwarding:

- A router must **NOT** forward a packet onto the subnet from which the packet was originated. This is achieved by performing a check on the source and the outgoing interface before building a source specific state or before forwarding a packet[18].

[18] Most implementations create a cache for forwarding packets. This check can be done only once when creating the cache, and is not done per packet.

Figure 11: The Join/Prune scenario under study

5 Related work

The related work falls mainly in the field of protocol verification. Most of the literature on multicast protocol design addresses architecture, specification, and comparisons between different protocols. We are not aware of any other work to develop systematic methods for testing multiparty protocol robustness. In addition, some concepts of STRESS were inspired by VLSI chip testing.

There is a large body of literature dealing with verification of communication protocols. Protocol verification typically addresses *safety*, *liveness*, and *responsiveness* properties [23]. Safety properties include freedom from deadlocks, assertion violations, improper terminations and unspecified receptions. Liveness properties include detection of acceptance cycles and absence of non-progress cycles, while responsiveness properties include timeliness and fault tolerance. Most protocol verification systems, including STRESS, aim to detect violations of (part of) these protocol properties.

In general, the two main approaches for protocol verification are theorem proving and reachability analysis (or model checking) [24, 25]. Theorem proving systems define a set of axioms and construct relations on these axioms. Desirable properties of the protocol are then proven mathematically. Theorem proving includes *model-based* formalisms (such as Z [26], and Vienna Development Method (VDM) [27]) and *logic-based* formalisms including first order logic (such as Nqthm [28]) and higher order logic (such as Prototype Verification System (PVS) [29]). Formal verification has been applied to TCP and T/TCP [30], but for multiparty protocols, theorem proving systems are likely to be even more complex and perhaps intractable.

Reachability analysis algorithms [1, 2] try to generate and inspect all the protocol states that are reachable from given initial state(s). Such algorithms suffer from the 'state space explosion' problem, espe-

This is different than the 'incoming interface' check stated by the current specification. In the specific case discussed here, the looping multicast packets will match on the incoming interface (the LAN) for the source-specific entry.

Figure 13: Generic BIST scheme

age with a minimum number of inputs. Similarly, the scenario filter chooses the representative scenarios to maximize protocol coverage with a minimum number of scenarios. Moreover, BIST uses a 'response monitor circuit' to monitor and detect error signals. This is analogous to our use of tracing and error detection modules. The expected output for VLSI chip testing is 'fault coverage vs. test length' curve, which is similar to 'protocol coverage vs. scenarios' statistics.

6 Summary and Future Work

The goals of our method are to simplify and systematize robustness analysis of multiparty protocols. Our approach is not about formal or abstract proofs. On the contrary, we attempt to provide practical meth-

ods to study robustness of real Internet multicast protocols, with the aid of systematic simulation[19]. This paper presented our initial attempts to achieve these goals in the context of one multicast routing protocol. We used scenario generation, simulation tracing, and output analysis to obtain a set of error-prone scenarios. In particular, we described several techniques:

- To circumvent the 'state explosion problem', we introduced the notion of *representative scenarios*. We obtained these scenarios for the multicast host model using a *scenario filter*, which excluded redundant and irrelevant scenarios based on practical assumptions.
- We identified two *representative topologies* based on the *equivalence* relationship established for the protocol under study, PIM-SM. The equivalence definition suggested that extending the simulated topologies would not reveal additional errors in the protocol.
- To capture robustness characteristics, we studied the protocol behavior in the presence of single packet loss. A *LAN* module accounted for the *selective loss* cases experienced by multicast messages. We hope to use a similar *logical LAN* module, to model loss and delay parameters of the underlying multicast distribution trees and thereby extend our method to higher-level multiparty protocols.
- To reduce the complexity of our analyses, we used *subsetting* of protocol functions, states and messages. This allowed abstracting out some protocol details while retaining and focusing on others.
- The definition of *error conditions*, in terms of end-point errors (such as data loss or duplication), enabled the output analyzer to capture *faulty scenarios* and isolate protocol traces in the recent history of the errors.
- Finally, we showed that a large portion of the protocol state space could be covered by simulating a few *representative scenarios*. With the aid of STRESS, we were able to discover several protocol design errors in PIM-SM, and suggest solutions to these errors.

This paper was the first attempt to develop and apply these methods. We are encouraged by our success in identifying protocol errors using these methods, and hope that we have similar results as we move on to investigate other multicast routing and end-to-end multiparty protocols.

Future directions for this research include:

- Developing algorithms for automatic scenario generation, that replace the heuristics used in this study. We are currently investigating semi-formal approaches, based on well-established VLSI testing techniques [39], and extending them to synthesize test topologies. In classical VLSI, test generation is performed on a given circuit. By contrast, network protocol robustness must be established over arbitrary and time-varying topologies.
- Further applying STRESS to multicast routing to investigate a wider range of protocol functions:
  - timers and timed actions, to complement the triggered actions investigated in this study.
  - heterogeneous topologies including asymmetric and uni-directional links, such as satellite links.

[19] We do not claim that the method presented here –based on heuristic scenario generation– is formal or general enough yet, and we are addressing these issues in our on-going and future work.

  - other multicast routing protocols, such as DVMRP, PIM-DM and hierarchical PIM [40].
  - interoperability between routing protocols. Examples include interaction between unicast and multicast routing, and the interoperability between multicast routing protocols.

- Generalizing the method and extending it to apply to end-to-end multiparty protocols. Examples of such protocols include reliable and real-time multicast transport and session management protocols, such as SRM, RTCP, and sdr. To achieve this, the multicast distribution tree may be viewed as a *logical LAN*, with various selective loss and delay models.

  Sensitivity analysis and performance evaluation may be needed for these protocols. For example, retransmission and congestion control mechanisms employed by end-to-end transport protocols usually use timer parameters. Methods may be needed for investigating the parameter space of the protocol (e.g. timer values) and the network (e.g. delay values), and study the effect of the change on performance in a systematic and efficient fashion.

- Applying to real implementation conformance testing [41] through an emulation interface. This facilitates driving the test using representative scenarios.

- Deriving behavioral assertion checks [42] based on the STRESS method. Assertions can be used in network management and self-diagnosing protocols.

7 Acknowledgements

We would like to thank Joseph Bannister (ISI), George Eddy (ISI), Sally Floyd (LBNL), Sandeep Gupta (USC), John Heideman (ISI), Kanna Kumar (USC/ISI), and Pavlin Radoslavov (USC/ISI) for their useful comments and insightful feedback on the paper.

References

[1] F. Lin, P. Chu, and M. Liu. Protocol Verification using Reachability Analysis. *Computer Communication Review, Vol. 17, No. 5*, 1987.

[2] F. Lin, P. Chu, and M. Liu. Protocol Verification using Reachability Analysis: the state explosion problem and relief strategies. *Proceedings of the ACM SIGCOMM'87*, 1987.

[3] D. Waitzman S. Deering, C. Partridge. Distance Vector Multicast Routing Protocol, November 1988. RFC1075.

[4] J. Moy. Multicast Extension to OSPF. *Internet Draft*, September 1992.

[5] D. Estrin, D. Farinacci, A. Helmy, V. Jacobson, and L. Wei. Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification. *Proposed Experimental RFC.*, September 1996.

[6] A. J. Ballardie, P. F. Francis, and J. Crowcroft. Core Based Trees. In *Proceedings of the ACM SIGCOMM*, San Francisco, 1993.

[7] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei. Protocol Independent Multicast - Sparse Mode (PIM-SM): Motivation and Architecture. *Proposed Experimental RFC.*, October 1996.

[8] S. Floyd, V. Jacobson, C. Liu, S. McCanne, and L. Zhang. A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing. *IEEE/ACM Transactions on Networking*, November 1996.

[9] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A Transport Protocol for Real-Time Applications. *RFC 1889*, January 1996.

[10] S. McCanne. A Distributed Whiteboard for Network Conferencing. *U.C. Berkeley Computer Science project*, May 1992.

[11] V. Jacobson and S. McCanne. vat - LBNL Audio Conferencing Tool. *URL www-nrg.ee.lbl.gov/vat*, 1993.

[12] S. McCanne and V. Jacobson. vic: A Flexible Framework for Packet Video. *ACM Multimedia 1995*, November 1995.

[13] M. Handley. NTE - The UCL Network Text Editor. 1996.

[14] M. Handley. The sdr Session Directory: An Mbone Conference Scheduling and Booking System. 1996.

[15] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei. Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification. *RFC 2117.*, March 1997.

[16] D. Estrin, M. Handley, A. Helmy, P. Huang, and D. Thaler. A Dynamic Bootstrap Mechanism for Rendezvous-based Multicast Routing. *Submitted to IEEE/ACM Transactions on Networking.*, May 1997.

[17] E. W. Dijkstra. A note on two problems in connection with graphs. *Numerische Mathematik, Vol. 1*, 1959.

[18] S. McCanne and S. Floyd. NS: Network Simulator. *URL www-nrg.ee.lbl.gov/ns*, 1995.

[19] J. Ousterhout. Tcl and the Tk Toolkit. *Addison Wesley*, 1994.

[20] D. Wetherall and C. Lindblad. Extending Tcl for Dynamic Object-Oriented Programming. *Proceedings of the Tcl/Tk Workshop 95, Toronto, Ontario*, July 1995.

[21] A. Helmy. Protocol Independent Multicast-Sparse Mode (PIM-SM): Implementation Document. *Internet Draft.*, January 1997.

[22] S. Floyd and V. Jacobson. Random Early Detection Gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking, Vol. 1, No. 4.*, pages 397–413, August 1993.

[23] K. Saleh, I. Ahmed, K. Al-Saqabi, and A. Agarwal. A recovery approach to the design of stabilizing communication protocols. *Journal of Computer Communication, Vol. 18, No. 4*, pages 276–287, April 1995.

[24] E. Clarke and J. Wing. Formal Methods: State of the Art and Future Directions. *ACM Workshop on Strategic Directions in Computing Research, Vol. 28, No. 4*, pages 626–643, December 1996.

[25] A. Helmy. A Survey on Kernel Specification and Verification. *TR 97-654 of CS at USC.*

[26] J. Spivey. Understanding Z: a Specification Language and its Formal Semantics. *Cambridge University Press*, 1988.

[27] C. Jones. Systematic Software Development using VDM. *Prentice-Hall Int'l*, 1990.

[28] R. Boyer and J. Moore. A Computational Logic Handbook. *Academic Press, Boston*, 1988.

[29] S. Owre, J. Rushby, N. Shanker, and F. Henke. Formal verification for fault-tolerant architectures: Prolegomena to the design of PVS. *IEEE Transactions on Software Engineering*, pages 107–125, February 1995.

[30] M. Smith. Formal Verification of Communication Protocols. *FORTE/PSTV'96 Conference*, October 1996.

[31] D. Probst. Using partial-order semantics to avoid the state explosion problem in asynchronous systems. *Proc. 2nd Workshop on Computer-Aided Verification, Springer Verlag, New York*, 1990.

[32] P. Godefroid. Using partial orders to improve automatic verification methods. *Proc. 2nd Workshop on Computer-Aided Verification, Springer Verlag, New York*, 1990.

[33] N. Maxemchuck and K. Sabnani. Probabilistic verification of communication protocols. *Proc. 7th IFIP WG 6.1 Int. Workshop on Protocol Specification, Testing, and Verification, North-Holland Publ., Amsterdam*, 1987.

[34] C. West. Protocol Validation by Random State Exploration. *Proc. 6th IFIP WG 6.1 Int. Workshop on Protocol Specification, Testing, and Verification, North-Holland Publ., Amsterdam*, 1986.

[35] J. Pageot and C. Jard. Experience in guiding simulation. *Proc. VIII-th Workshop on Protocol Specification, Testing, and Verification, Atlantic City, 1988, North-Holland Publ., Amsterdam*, 1988.

[36] G. Holzmann. Design and Validation of Computer Protocols. *AT&T Bell Labs., Prentice Hall*, 1991.

[37] B. Murray and J. Hayes. Testing ICs: Getting to the Core of the Problem. *IEEE Computer Magazine*, pages 32–38, November 1996.

[38] B. Konemann, B. Bennetts, N. Jarwala, and B. Nadeau-Dostie. Built-In Self-Test: Assuring System Integrity. *IEEE Computer Magazine*, pages 39–45, November 1996.

[39] M. Abramovici, M. Breuer, and A. Friedman. Digital Systems Testing and Testable Design. *AT & T Labs.*, 1990.

[40] S. Deering, B. Fenner, D. Estrin, A. Helmy, D. Farinacci, L. Wei, M. Handley, V. Jacobson, and D. Thaler. Hierarchical PIM-SM Architecture for Inter-Domain Multicast Routing. *Internet Draft.*, December 1995.

[41] D. Rayner. OSI conformance testing. *Computer Networks and ISDN Systems, Special issue on Conformance Testing, Vol. 14, No. 1*, pages 79–98, 1987.

[42] S. Perl. Performance Assertion Checking. *Ph.D. Thesis, MIT*, September 1992.