

Byzantine Replication Under Attack *

Yair Amir¹, Brian Coan², Jonathan Kirsch¹, John Lane¹

¹ Johns Hopkins University, Baltimore, MD. {yairamir, jak, johnlane}@cs.jhu.edu

² Telcordia Technologies, Piscataway, NJ. coan@research.telcordia.com

Abstract

Existing Byzantine-resilient replication protocols satisfy two standard correctness criteria, safety and liveness, in the presence of Byzantine faults. In practice, however, faulty processors can, in some protocols, significantly degrade performance by causing the system to make progress at an extremely slow rate. While “correct” in the traditional sense, systems vulnerable to such performance degradation are of limited practical use in adversarial environments.

This paper argues that techniques for mitigating such performance attacks are needed to bridge this “practicality gap” for intrusion-tolerant replication systems. We propose a new performance-oriented correctness criterion, and we show how failure to meet this criterion can lead to performance degradation. We present a new Byzantine replication protocol that achieves the criterion and evaluate its performance in fault-free configurations and when under attack.

Keywords: Byzantine, replication, fault tolerance, performance, attacks

1 Introduction

Existing Byzantine-resilient state machine replication (SMR) protocols satisfy two standard correctness criteria in the presence of Byzantine faults: safety and liveness. Safety means that two servers remain consistent replicas of one another, while liveness means that each update is executed eventually. Since no asynchronous Byzantine agreement protocol can always be both safe and live [14], systems requiring strong consistency semantics are usually designed to meet safety in all executions, while guaranteeing liveness only during periods of sufficient synchrony and connectivity [13] or in a probabilistic sense [5, 23].

Designers of practical Byzantine-resilient replication systems recognize that real systems are not completely

asynchronous. Rather, these systems exhibit extended periods of stability (synchrony), possibly interspersed with periods of instability. Realistic Byzantine-resilient replication systems generally guarantee liveness in a sufficiently stable subset of the set of all asynchronous executions. In this paper we observe that during stable periods, the system can satisfy much stronger performance guarantees. Thus, when the network is stable, there is a potential gap in the type of performance that is promised by existing protocols (i.e., eventual execution of each update) and the type of performance that is attainable.

In Byzantine environments, faulty processors can exploit this gap to degrade system performance to a level far below what would be achievable with only correct processors. Specifically, a small number of faulty processors can cause the system to make progress at an extremely slow rate. While “correct” in the traditional sense (both safety and liveness are met), systems vulnerable to such performance degradation are of limited practical use in adversarial environments.

We experienced this problem first hand during a red-team experiment conducted on our Steward system [4]. Although the system survived all of the tests according to the metrics of safety and liveness, we observed that it was slowed down to twenty percent of its potential performance in one experiment. After analyzing the attack, we found that we could in fact slow the system down to roughly one percent of its potential performance. Thus, our provably correct system, which achieves high performance in fault-free configurations, could be made effectively unusable in practice under a relatively simple attack. This experience led us to conclude that liveness is a necessary but insufficient correctness criterion for achieving high performance Byzantine replication under attack. This paper argues that new *performance-oriented* criteria are needed.

Preventing the type of performance degradation experienced by Steward requires addressing what we call a *Byzantine performance failure*. Previous work focused on Byzantine failures in the value domain (where a faulty processor tries to subvert the protocol by sending incorrect or conflicting messages) and the time domain (where messages

*This publication was supported by grants 0430271 and 0716620 from the National Science Foundation. Its contents are solely the responsibility of the authors and do not necessarily represent the official view of Johns Hopkins University or the National Science Foundation.

from a faulty processor do not arrive within protocol timeouts, if at all). Processors exhibiting performance failures, however, send correct messages slowly but without triggering protocol timeouts; they are thus correct in both of the traditional domains, despite having the potential to significantly degrade performance. Performance failures have been considered in benign environments [11, 25]. To the best of our knowledge, we are the first to (1) propose a useful performance-oriented metric to evaluate Byzantine protocols and (2) present a SMR protocol that performs well according to this metric.

Byzantine protocols whose progress is driven by messages from a large number of correct processors (e.g., [5, 22]) are less vulnerable to performance degradation due to performance failures. The voting in such protocols masks performance failures, in addition to value and timing failures, because no collection of faulty processors can prevent the correct processors from moving forward. For efficiency, however, other protocols rely on select processors to perform certain tasks correctly and in a timely manner, reducing the number of messages that must be sent in the common case. These protocols typically use cryptographic tools and timeouts to restrict the adversary in the value and time domains, respectively, but they do not address performance failures.

In this paper we focus on this latter class of Byzantine SMR protocols, which we refer to as *leader-based* protocols. These protocols (e.g., [2, 4, 8, 16, 18, 21, 29]) rely on a leader to coordinate the global ordering and are thus vulnerable to performance degradation caused by a slow leader. The problem is magnified in environments (such as wide-area networks) where it is difficult to predict the type of performance that should be expected of the leader. We demonstrate this vulnerability through analysis and experimental evaluation of BFT [8], the first leader-based Byzantine fault-tolerant SMR protocol to achieve practical performance in fault-free executions.

By applying the understanding gained from our experience with BFT, we developed a new Byzantine fault-tolerant SMR protocol, Prime (Performance-oriented Replication In Malicious Environments), resilient to performance degradation under attack. Prime has two key properties: (1) The resources required by the leader for global ordering are bounded and independent of system throughput, enabling non-leader servers to aggressively monitor the leader’s performance, and (2) Non-leader servers compute a threshold level of acceptable performance, which is a function of current network latencies, against which they judge the leader. Prime meets a new performance-oriented correctness criterion, BOUNDED-DELAY, which makes a stronger guarantee than traditional liveness criteria. We present experimental results showing that Prime performs competitively with BFT in fault-free configurations and performs an order of

magnitude better when under attack. Our results show that the performance of Prime when under attack is within a reasonable factor of its fault-free performance.

The remainder of this paper is presented as follows. Section 2 presents our system model and describes the service properties provided by our system. Section 3 describes the vulnerabilities of existing leader-based protocols to performance degradation under attack, using BFT as a case study. We present the Prime protocol in Section 4, and we analyze its properties in Section 5. Section 6 presents experimental results for our new system. Section 7 details related work, and Section 8 concludes the paper.

2 System Model and Service Properties

We consider a system consisting of N servers, which communicate by passing messages. Each server is uniquely identified from the set $\mathcal{R} = \{1, 2, \dots, N\}$. We assume a Byzantine fault model. Servers are either *correct* or *faulty*; correct servers follow the protocol specification, while faulty servers can deviate from the protocol specification arbitrarily. We employ digital signatures, and we make use of a cryptographic hash function to compute message digests. We denote a message m signed by server i as $\langle m \rangle_{\sigma_i}$, and we denote a digest of m as $D(m)$. We assume that all adversaries, including faulty servers, are computationally bounded such that they cannot subvert these cryptographic mechanisms.

The consistency of our new protocol, Prime, is given in the following two properties:

DEFINITION 2.1 SAFETY: *If two correct servers execute the i^{th} update, then these updates are identical.*

DEFINITION 2.2 VALIDITY: *Only an update that was proposed by a client may be executed.*

Prime guarantees safety and validity in all executions, including those in which the network is asynchronous and may drop or duplicate messages. Like existing leader-based Byzantine replication protocols, Prime guarantees liveness only in executions in which the network eventually meets certain stability conditions, which we now state. In what follows, K_{Lat} is a known network-specific constant accounting for latency variability.

DEFINITION 2.3 PRIME-STABILITY: *There is a time after which the following condition holds for a set of at least $2f + 1$ correct servers (the stable servers):*

- *For each pair of stable servers r and s , there exists a value $Min_Lat(r, s)$, unknown to the servers, such that if r sends a message to s , it will arrive with delay $\Delta_{r,s}$, where $Min_Lat(r, s) \leq \Delta_{r,s} \leq Min_Lat(r, s) * K_{Lat}$.*

In those executions in which PRIME-STABILITY is met, Prime guarantees the following liveness property:

DEFINITION 2.4 PRIME-LIVENESS: *If a stable server initiates an update, all stable servers will eventually execute the update.*

PRIME-LIVENESS is similar to the liveness guarantees provided by existing leader-based protocols (except that PRIME-LIVENESS contains a stronger degree of stability). While it is critical to guarantee that in those executions that are sufficiently stable each update is eventually executed, such liveness properties do not guarantee how quickly the updates are executed when the network is stable. Systems that solely meet liveness thus provide a very weak performance-related guarantee.

For this reason, in those executions in which PRIME-STABILITY is met, Prime also provides a stronger performance guarantee, which we call BOUNDED-DELAY:

DEFINITION 2.5 BOUNDED-DELAY: *There exists a time after which the update latency for any update initiated by a stable server is upper-bounded.*

Prime achieves BOUNDED-DELAY in those executions in which PRIME-STABILITY is met, assuming the system is not overloaded (i.e., given load beyond its maximum throughput) and when correct servers have sufficient bandwidth with which to communicate. Indeed, no system (even in benign environments) can provide latency guarantees when these conditions are not met due to necessary queuing delays. Our current protocol requires knowledge of this minimal level of bandwidth to ensure that these assumptions are met. We believe that adaptively setting the bandwidth consumed by correct servers is an important open problem for Byzantine-resilient systems. Section 5 provides an analysis of the bound provided by Prime.

We remark that resource exhaustion denial of service attacks may cause PRIME-STABILITY to be violated for the duration of the attack. However, such attacks fundamentally differ from the attacks that are the focus of this paper, where malicious leaders can slow down the system without triggering defense mechanisms (see Section 3). Handling resource exhaustion attacks is a difficult problem that is orthogonal and complementary to the solution strategies considered in this paper.

3 Case Study: BFT Under Attack

In this section we present a theoretical analysis of BFT [8], a leader-based Byzantine SMR protocol, when under attack. We chose BFT because (1) it is the standard protocol to which other Byzantine protocols are often compared, (2) many of the attacks that can be applied to BFT (and the

corresponding lessons learned) also apply to other leader-based protocols, and (3) its implementation was publicly available. BFT achieves high throughputs in fault-free configurations or when servers exhibit only benign faults. We first provide background on BFT and then describe two attacks that can be used to significantly degrade its performance when under attack. We present experimental results validating the analysis in Section 6.

BFT assigns a total order to client updates. The protocol requires $3f + 1$ servers, where f is the maximum number of servers that may be Byzantine. An elected leader coordinates the protocol by assigning sequence numbers to updates. If a server suspects that the leader has failed, it votes to replace it. When $2f + 1$ servers vote to replace the leader, a view change occurs, in which a new leader is elected and servers collect information regarding pending updates so that progress can safely resume in a new view.

A client sends its updates directly to the leader. The leader assigns a sequence number to the update and proposes the assignment to the rest of the servers. It sends a PRE-PREPARE message, which contains the view number, the assigned sequence number, and the update itself. Upon receiving the PRE-PREPARE, a non-leader server accepts the proposed assignment by broadcasting a PREPARE message. The PREPARE message contains the view number, the assigned sequence number, and a digest of the update. When a server collects the PRE-PREPARE and $2f$ corresponding PREPARE messages, it broadcasts a COMMIT message. A server globally orders the update when it collects $2f + 1$ COMMIT messages. Each server executes globally ordered updates according to sequence number. A server sends a reply to the client after executing the update.

3.1 Attack 1: Pre-Prepare Delay

A malicious leader can introduce latency into the global ordering path simply by waiting some amount of time after receiving an update before sending it in a PRE-PREPARE message. The amount of delay a leader can add without being detected as faulty is dependent on (1) the way in which non-leaders place timeouts on updates they have not yet executed and (2) the duration of these timeouts.

A malicious leader can ignore updates sent directly by clients. If a client's timeout expires before receiving a reply to its update, it broadcasts the update to all servers, which forward the update to the leader. Each non-leader server maintains a FIFO queue of pending updates (i.e., those updates it has forwarded to the leader but not yet executed). A server places a timeout on the execution of the first update in its queue; that is, it expects to execute the update within the timeout period. If the timeout expires, the server suspects the leader is faulty and votes to remove it from power. When a server executes the first update in the queue,

it restarts the timer if the queue is not empty. Note that a server does not stop the timer if it executes a pending update that is not the first in the queue. The duration of the timeout is dependent on its initial value (which is implementation and configuration dependent) and the history of past view changes. Servers double the value of their timeout each time a view change occurs. The specification of BFT does not provide a mechanism for reducing timeout values.

BFT’s queuing mechanism ensures fairness by guaranteeing that each update is eventually ordered. However, it also allows the leader to significantly delay the ordering of an update without being replaced. To stay in power, the leader must prevent $f + 1$ correct servers from voting to replace it. Thus, assuming a timeout value of TO , a malicious leader can use the following attack: (1) Choose a set S of $f + 1$ correct servers, (2) For each server $r \in S$, maintain a FIFO queue of the updates forwarded by r , and (3) For each such queue, send a PRE-PREPARE containing the first update on the queue every $TO - \epsilon$ time units. This guarantees that the $f + 1$ correct servers in S execute the first update on their queue each timeout period. If these updates are all different, the fastest the leader would need to introduce updates is at a rate of $f + 1$ per timeout period. In the worst case, the $f + 1$ servers would have identical queues, and the leader could introduce one update per timeout.

This attack exploits the fact that non-leader servers place timeouts only on the first update in their queues. To understand the ramifications of placing a timeout on *all* pending updates, consider the following scenario: Non-leader server s simultaneously initiates n updates. If server s sets a timeout on all n updates, then s will suspect the leader if the system fails to execute n updates per timeout period. Since the system has a maximal throughput, if n is sufficiently large, s will suspect a correct leader. The fundamental problem is that correct servers have no way to assess the rate at which a correct leader can coordinate global ordering.

3.2 Attack 2: Timeout Manipulation

One of the main benefits of BFT is that it ensures safety regardless of synchrony assumptions. The authors justify the need for this property by noting that denial of service attacks can be used by a malicious adversary to violate timing assumptions. While a DoS attack cannot impact safety, it can be used to increase the timeout value used to detect a faulty leader. During the attack, the timeout doubles with each view change. If the adversary stops the attack when a malicious leader is in power, then that leader will be able to slow the system down to a throughput of roughly $f + 1$ updates per TO , where TO is potentially very large, using the attack described in the previous section. This vulnerability stems from the inability of BFT to reduce the timeout and adapt to the network conditions after the system stabilizes.

4 The Prime Protocol

In this section we present Prime, a new Byzantine fault-tolerant state machine replication protocol designed to mitigate the types of attacks described in Section 3. Prime requires $3f + 1$ servers to tolerate f Byzantine faults.

4.1 Prime Ordering Protocol

Prime uses a rotating coordinator protocol to assign a total order to client updates. The servers execute the updates according to this total order, and they thus remain replicas of one another. Prime establishes the total order in two phases. In the first phase, each server disseminates its updates to the other servers and coordinates an agreement protocol, which *preorders* those updates that it originated. Each preordering agreement protocol coordinated by a different server operates independently and in parallel. A preordered update, u , is bound to a *preorder identifier*, (o, i) , where u is the i th update preordered by server o . Thus, the preordering phase enables correct servers to consistently refer to updates using their preorder identifiers. In the second phase, an elected leader coordinates a global ordering protocol, which establishes a total order on batches of preordered updates. The final total order on updates is achieved by deterministically assigning an order to the updates in each batch based on their preorder identifiers.

Preordering Phase: When originating server o receives update u from one of its clients, it sends a $\langle \text{PO-REQUEST}, seq, u, o \rangle_{\sigma_o}$ message, req , to all servers, where seq is a local sequence number that o increments each time it sends a new PO-REQUEST. We refer to this local sequence number as a *preorder sequence number*. Upon receiving req , each correct server, i , sends a $\langle \text{PO-ACK}, seq, D(u), o, i \rangle_{\sigma_i}$ message to all other servers if i has not previously received a PO-REQUEST from o with sequence number seq . A set consisting of req and $2f$ matching PO-ACK messages constitutes a *preorder-certificate*, which is proof that the correct servers agree that preorder identifier (o, seq) is uniquely bound to u .

Each server, i , maintains a vector, $\text{PO_Aru}[]$, where $\text{PO_Aru}[o]$ contains the maximum sequence number, n , such that i has preorder-certificates for all preordered updates with identifiers (o, j) , with $j \leq n$. Each server, i , periodically broadcasts a $\langle \text{PO-ARU}, vec, i \rangle_{\sigma_i}$ message, where vec is its local PO_Aru vector. The PO-ARU message serves as a cumulative acknowledgement for preordered updates. Given two PO-ARU messages, m_1 and m_2 , Figure 1 defines what it means for m_1 to be *at least as up-to-date as* m_2 , *more up-to-date than* m_2 , and *consistent with* m_2 . Each server stores the most up-to-date, consistent PO-ARU message received from each other server in a vector, $\text{Last_PO_Aru}[]$, indexed by server identifier. We describe a

1. For $m_1 = \langle \text{PO-ARU}, \text{vec}_1, i \rangle_{\sigma_i}$ and $m_2 = \langle \text{PO-ARU}, \text{vec}_2, i \rangle_{\sigma_i}$, we say that:
 - m_1 is at least as up-to-date as m_2 when $(\forall j \in \mathcal{R})[\text{vec}_1[j] \geq \text{vec}_2[j]]$.
 - m_1 is more up-to-date than m_2 when m_1 is at least as up to date as $m_2 \wedge (\exists j \in \mathcal{R})[\text{vec}_1[j] > \text{vec}_2[j]]$.
 - m_1 and m_2 are consistent when m_1 is at least as up to date as m_2 , or m_2 is at least as up to date as m_1 .
2. For originating server o and preorder sequence number po_seq , $\text{Preorder_Proof_Exists}(o, po_seq, \langle \text{PRE-PREPARE}, *, *, \text{pm}, l \rangle_{\sigma_l})$ is true iff:
 - $|\{i : i \in \mathcal{R} \wedge \text{pm}[i][o] \geq po_seq\}| \geq 2f + 1$
3. $M(pp) = \langle \text{PRE-PREPARE}, *, *, \text{seq}, *, * \rangle_{\sigma_*} = \{(o, s) : o \in \mathcal{R} \wedge s \in \mathbb{N} \wedge \text{Preorder_Proof_Exists}(o, s, pp)\}$
4. For $pp = \langle \text{PRE-PREPARE}, *, *, p_{pp}, * \rangle_{\sigma_*}$ and $pm = \langle \text{PROOF-MATRIX}, p_{pm}, * \rangle_{\sigma_*}$, where p_{pp} and p_{pm} denote proof matrices, we say that:
 - pp covers pm if $\forall i \in \mathcal{R}, p_{pp}[i]$ is at least as up-to-date as $p_{pm}[i]$.

Figure 1: Definitions and terminology used by the Prime ordering protocol.

mechanism for blacklisting faulty servers that send inconsistent PO-ARU messages in [3].

Global Ordering Phase: Prime’s global ordering phase is similar to BFT and uses three message rounds (see Section 3). While BFT establishes a total order on PRE-PREPARE messages containing updates, Prime’s global ordering phase establishes a total order on PRE-PREPARE messages containing *proof matrices*. Each proof matrix is a vector of PO-ARU messages. A correct leader, l , periodically sends a $\langle \text{PRE-PREPARE}, v, \text{seq}, \text{pm}, l \rangle_{\sigma_l}$ message, where v is the current view number, seq is a global sequence number, and pm is the leader’s Last_PO_Aru vector (which is a proof matrix). $\text{pm}[o]$ is either a PO-ARU message signed by server o or a null vector of length $|\mathcal{R}|$, indicating that o has not yet cumulatively acknowledged any preorder-certificates.

We now explain how a server obtains a total order on updates from the totally ordered stream of PRE-PREPARE messages. Call this stream of PRE-PREPARE messages $T = \langle T_1, T_2, \dots \rangle$. Intuitively, globally ordering a PRE-PREPARE message expands the set of preordered updates that are eligible for execution. Let M map a globally ordered PRE-PREPARE, pp , to a set of preordered updates, P , where P contains those preordered updates, (o, s) , for which $\text{Preorder_Proof_Exists}(o, s, pp)$ is true (see Figure 1). Let L be a function that lexicographically orders the elements of P by their preorder identifiers. Then the final total order, U , on updates is obtained by $U = L(M(T_1)) \parallel L(M(T_2) - M(T_1)) \parallel L(M(T_3) - M(T_2)) \dots$, where \parallel denotes concatenation and $-$ denotes set difference.

Part A of Figure 2 summarizes the path of an update, u , through the system in the fault-free case. The update is pre-




Figure 2: Common case operation of Prime ($f = 1$). Part A shows the messages and protocol rounds when the leader is correct. Part B shows the delay added by a malicious leader that performs well enough to stay in power. The malicious leader ignores PO-ARU messages and sends its PRE-PREPARE to only one correct server.

ordered in two rounds, after which its preordering is cumulatively acknowledged in PO-ARU messages. When the network is stable, faulty servers cannot delay the preordering of u because correct servers need only wait for PO-ACK messages from each other to collect a preorder-certificate for u . In turn, the faulty servers cannot delay how quickly the preordering of u is cumulatively acknowledged in the PO-ARU messages of correct servers. A correct leader sends a PRE-PREPARE, pp , whose proof matrix includes these PO-ARU messages. u will be executed when pp is globally ordered.

Reconciliation: In order to prevent attacks where malicious servers preorder updates without sending them to all correct servers, Prime includes a reconciliation mechanism that disseminates updates to correct servers that do not receive them during the preordering phase. Together, Prime’s preordering phase and its reconciliation protocol provide a reliable broadcast service; if Prime assigns a global order to update u , reconciliation guarantees that all correct servers will receive u . Reconciliation messages contain $\text{MDS}(2f+1, f+1)$ erasure encoded [6] PO-REQUESTS, such that a server needs to receive $f+1$ out of $2f+1$ reconciliation messages to decode the associated PO-REQUEST. This reduces the bandwidth overhead required for reconciliation, and is described in detail in [3].

4.2 Detecting Malicious Leaders

A malicious leader can mount two types of performance attacks against Prime. First, it can propose a global ordering on preordered updates slowly by sending PRE-PREPARE messages at a slow rate. Some strategies for the leader to slow down the sending of its PRE-PREPARES are illustrated in Part B of Figure 2. Prime uses the SUSPECT-LEADER protocol, described below, to detect slow leaders. Second, even if it sends timely PRE-PREPARE messages, a malicious leader can intentionally send a PRE-PREPARE, pp , whose proof matrix does not contain the most up-to-date PO-ARU messages that it has received. This can prevent preordered updates that would have become eligible for execution when pp is globally ordered from becoming eligi-

ble. Defending against these two performance attacks allows Prime to meet BOUNDED-DELAY (see Definition 2.5).

Enforcing up-to-date Pre-Prepare messages: Each non-leader server, i , periodically sends a $\langle \text{PROOF-MATRIX}, pm, i \rangle_{\sigma_i}$ message to the leader, where pm is i 's $\text{Last_PO_Aru}[]$. pm contains the set of PO_Aru vectors i expects the leader to include in its next PRE-PREPARE . To understand why a non-leader server is justified in this expectation, note that the leader can simply adopt any of the PO_Aru vectors in pm that are more up-to-date than what it currently has in its $\text{Last_PO_Aru}[]$. Thus, a correct leader will send, in its next PRE-PREPARE , a proof matrix with vectors that are at least as up-to-date as those in pm . We say that such a PRE-PREPARE covers pm (see Figure 1). A critical property of Prime, which differs from existing leader-based solutions, is that the leader requires a bounded amount of bandwidth and computational resources, independent of system throughput, to perform its role as leader; the size of a PRE-PREPARE is dependent only on the number of servers, and a single PRE-PREPARE can propose a global ordering on an arbitrary number of preordered updates.

Pre-Prepare Flooding: Prime's mechanism for detecting malicious leaders requires a simple addition to the global ordering phase to ensure timely global ordering. Upon receiving a PRE-PREPARE , pp , a correct server broadcasts it. This guarantees that all correct servers receive pp within one round from the time that the first correct server receives it, at which point no faulty server can delay the correct servers from globally ordering pp . Flooding PRE-PREPARE s forces a malicious leader to delay sending PRE-PREPARE s to all correct servers in order to add unbounded delay to the global ordering phase. In practice, the rate at which the leader sends PRE-PREPARE s can be configured so that this flooding requires a small bandwidth overhead.

Suspect-Leader Protocol: Since the leader requires bounded resources to perform its role as leader, if the network is stable, the leader can be expected to send up-to-date PRE-PREPARE s in a timely manner. To leverage this, we require a mechanism whereby non-leader servers can (1) dynamically determine how fast a timely leader should perform, (2) monitor the performance of the current leader, and (3) suspect the leader if it is not performing fast enough. Each time a server sends a PROOF-MATRIX message, pm , it computes the delay between sending pm and receiving a PRE-PREPARE covering pm . We call this delay the *turn-around-time* (abbreviated TAT) provided by the leader. The goal of SUSPECT-LEADER is to force any leader that stays in power to provide a timely TAT to at least one correct server.

Each correct server, i , locally decides whether to suspect the leader by computing two values, TAT_acceptable and TAT_leader . TAT_acceptable is a standard against which server i judges the current leader, and TAT_leader is a measure of the current leader's performance. Server i suspects

the leader if $\text{TAT_leader} > \text{TAT_acceptable}$.

TAT_acceptable and TAT_leader are computed so that, when PRIME-STABILITY holds, SUSPECT-LEADER meets two key properties. L^* is the maximum latency between any two correct servers after the network stabilizes. Δ_{pp} is a value greater than the maximum time between a correct server sending successive PRE-PREPARE messages. K_{Lat} (see Section 2) accounts for latency variability. Let $B = 2K_{Lat}L^* + \Delta_{pp}$. We now state the properties.

PROPERTY 4.1 *Any server that retains a role as leader must provide a turn-around time to at least one correct server that is no more than B .*

PROPERTY 4.2 *There exists a set of at least $f + 1$ correct servers (the permanent leaders) that will not be suspected by any correct server if elected leader.*

Intuitively, Property 4.1 ensures that a faulty leader will be suspected unless it provides a timely TAT to at least one correct server. We consider a TAT, $t \leq B$, to be timely because B is within a constant factor of the TAT that the slowest correct server might provide. This factor is a function of the latency variability that SUSPECT-LEADER is configured to tolerate. Note that malicious servers cannot affect the value of B . Property 4.2 ensures that view changes cannot occur indefinitely. Prime does not guarantee that the slowest f correct servers will not be suspected because slow faulty leaders cannot be distinguished from slow correct leaders.

Figure 3 contains pseudocode for SUSPECT-LEADER . Server i initializes its data structures at the beginning of each new view (Block A). The remaining blocks run in parallel. In Block B, server i uses a simple ping protocol to measure the RTT to each other server, j . Server i sends this measured RTT to j . Using this value, j computes the maximum TAT that i would compute for j if j were the leader, and stores it in $\text{TATs_If_Leader}[i]$. In Block C, server i uses $\text{TATs_If_Leader}[]$ to compute an upper bound, α , on the value of TAT_Leader that any correct server will compute for i if it were leader. Each server broadcasts its value of α and stores the values that it receives in $\text{TAT_Leader_UBs}[]$. In Block D, each non-leader server broadcasts the maximum TAT that the leader has provided it in the current view and stores the values that it receives in $\text{Reported_TATs}[]$. In Block E, each server computes TAT_acceptable using $\text{TAT_Leader_UBs}[]$, computes TAT_Leader using $\text{Reported_TATs}[]$, and compares these values to decide whether to suspect the leader.

We now sketch the proof of Property 4.1. From Block B, at least $2f + 1$ cells in server i 's $\text{TATs_If_Leader}[]$ eventually contain values, v , sent by correct servers. By definition, each $v \leq B$. Figure 4 (left side) illustrates that server i computes an α such that $\alpha \leq B$ (see Line C3). Server i stores the values of α computed by each other server. Thus,

```

/* Initialization, run at start of new view */
A1. For  $i = 1$  to  $N$ ,  $TATs\_If\_Leader[i] \leftarrow \infty$ 
A2. For  $i = 1$  to  $N$ ,  $TAT\_Leader\_UBs[i] \leftarrow \infty$ 
A3. For  $i = 1$  to  $N$ ,  $Reported\_TATs[i] \leftarrow 0$ 
A4.  $ping\_seq \leftarrow 0$ 

/* RTT Measurement Task, run at server  $i$  */
B1. Periodically:
B2. BROADCAST:  $\langle RTT-PING, view, ping\_seq, i \rangle_{\sigma_i}$ 
B3.  $ping\_seq++$ 
B4. Upon receiving  $\langle RTT-PING, view, seq, j \rangle_{\sigma_j}$ :
B5. SEND to server  $j$ :  $\langle RTT-PONG, view, seq, i \rangle_{\sigma_i}$ 
B6. Upon receiving  $\langle RTT-PONG, view, seq, j \rangle_{\sigma_j}$ :
B7.  $rtt \leftarrow$  Measured RTT for pong message
B8. SEND to server  $j$ :  $\langle RTT-MEASURE, view, rtt, i \rangle_{\sigma_i}$ 
B9. Upon receiving  $\langle RTT-MEASURE, view, rtt, j \rangle_{\sigma_j}$ :
B10.  $t \leftarrow rtt * K_{Lat} + \Delta_{pp}$ 
B11. if  $t < TATs\_If\_Leader[j]$ 
B12.  $TATs\_If\_Leader[j] \leftarrow t$ 

/* TATLeader Upper Bound Task, run at server  $i$  */
C1. Periodically:
C2.  $Sorted\_TATs \leftarrow$  SORT-ASCENDING  $TATs\_If\_Leader[]$ 
C3.  $\alpha \leftarrow Sorted\_TATs[2f+1]$ 
C4. BROADCAST:  $\langle TAT-UB, view, \alpha, i \rangle_{\sigma_i}$ 
C5. Upon receiving  $\langle TAT-UB, view, tat\_ub, j \rangle_{\sigma_j}$ :
C6. if  $tat\_ub < TAT\_Leader\_UBs[j]$ 
C7.  $TAT\_Leader\_UBs[j] \leftarrow tat\_ub$ 

/* TAT Measurement Task, run at server  $i$  */
D1. Periodically:
D2.  $max\_tat \leftarrow$  Maximum TAT measured this view
D3. BROADCAST:  $\langle TAT-MEASURE, view, max\_tat, i \rangle_{\sigma_i}$ 
D4. Upon receiving  $\langle TAT-MEASURE, view, tat, j \rangle_{\sigma_j}$ 
D5. if  $tat > Reported\_TATs[j]$ 
D6.  $Reported\_TATs[j] \leftarrow tat$ 

/* Suspect Leader Task */
E1. Periodically:
E2.  $Sorted\_TAT\_UBs \leftarrow$  SORT-ASCENDING  $TAT\_Leader\_UBs[]$ 
E3.  $TAT\_acceptable \leftarrow Sorted\_TAT\_UBs[2f+1]$ 
E4.  $Sorted\_TATs \leftarrow$  SORT-ASCENDING  $Reported\_TATs[]$ 
E5.  $TAT\_leader \leftarrow Sorted\_TATs[f+1]$ 
E6. if  $TAT\_leader > TAT\_acceptable$ 
E7. Suspect Leader

```

Figure 3: SUSPECT-LEADER Protocol, used to determine whether a server should suspect the leader. View numbers refer to the view in the global ordering protocol.

at least $2f + 1$ of the cells in server i 's $TAT_Leader_UBs[]$ (Figure 4, right side) eventually contain α values from correct servers (each of which is no more than B). The right side of Figure 4 shows that the value of $TAT_acceptable$ (denoted τ) must also eventually be less than or equal to B .

If a malicious leader remains in power, there are at least $f + 1$ servers (at least one of which is correct) for which $TAT_leader \leq TAT_acceptable$ always holds. Thus, at least one correct server collects $TAT-MEASURE$ messages from $f + 1$ servers (at least one of which is correct) with values v such that $v \leq TAT_acceptable$. Therefore, the malicious leader is providing a TAT , t , such that $t \leq TAT_Acceptable \leq B$, to at least one correct server.

We now show that Property 4.2 holds. Since $TAT_acceptable$ is the $(2f + 1)$ st lowest value in $TAT_Leader_UBs[]$, at least $f + 1$ correct servers sent values for α such that $\alpha \leq TAT_acceptable$. Each permanent leader, l , has a set of at least $f + 1$ correct servers that, if l is elected, will report $TATs$, t , with $t \leq \alpha \leq TAT_acceptable$. Thus, any correct server will compute $TAT_leader \leq TAT_acceptable$ and will not suspect l .




Figure 4: The value of $TAT_acceptable$ (τ) computed at any correct server converges to a value where $\tau \leq B$. The sets M (left side) and P (right side) contain the $f + 1$ highest values in their respective vectors. M must eventually contain at least one value, v , reported by a correct server, where $v \leq B$. Thus, $\alpha \leq B$. Using a parallel argument, the right side shows that $\tau \leq B$.

5 Analysis

In this section we show that in those executions in which PRIME-STABILITY holds, Prime meets the BOUNDED-DELAY property (see Definition 2.5). L^* and B are as defined in Section 4.2. Δ_{agg} is a value greater than the maximum time between a correct server sending any of the following messages successively: PO-ARU, PROOF-MATRIX, and PRE-PREPARE.

We first consider the maximum amount of delay that can be added by a malicious leader that performs well enough to stay in power. As discussed in Section 4, the time between a correct server initiating an update, u , and all correct servers sending PROOF-MATRIX messages containing at least $2f + 1$ PO-ARUs that cumulatively acknowledge the preordering of u is at most three rounds plus $2\Delta_{agg}$. The malicious servers cannot increase this time beyond what it would take if only correct servers were participating. By Property 4.1, a leader that stays in power must provide a TAT , $t \leq B$, to at least one correct server. By definition, $\Delta_{agg} \geq \Delta_{pp}$. Thus, $B \leq 2K_{Lat}L^* + \Delta_{agg}$. Since correct servers flood PRE-PREPARE messages, all correct servers receive the PRE-PREPARE within three rounds and one aggregation delay of when the PROOF-MATRIX messages are sent. All correct servers globally order the PRE-PREPARE in two rounds from the time, t , the last correct server receives it. Reconciliation guarantees that all correct servers receive the update within one round of time t . Summing the total delays yields a maximum latency of $\beta = 6L^* + 2K_{Lat}L^* + 3\Delta_{agg}$.

If a malicious leader delays an update by more than B , it will be suspected and a view change will occur. View changes require a finite (and, in practice, small) amount of state to be exchanged among correct servers, and thus they complete in finite time. SUSPECT-LEADER guarantees that at most $2f$ view changes can occur before the system settles on a leader that will remain in power forever. Therefore, there is a time after which the bound of β holds for any update initiated by a stable server.




Figure 5: Throughput (updates/sec) vs. number of clients, 50 ms Diameter

6 Performance Evaluation

To evaluate the performance of Prime, we implemented the protocol and compared it to an available implementation of BFT in a 7 server configuration ($f = 2$). We present results for $f = 1$ in [3]. Prime has similar performance to BFT when the systems are run in a benign environment, which is commonly the only environment in which Byzantine fault-tolerant replication systems are benchmarked. When strong attacks are mounted against both systems, Prime outperforms BFT by more than an order of magnitude.

Testbed and Network Setup: We used a system consisting of 7 servers, organized in a fully connected graph. Each server ran on a 3.2 GHz, 64-bit Intel Xeon computer. RSA signatures provided authentication and non-repudiation. Each computer can compute a 1024-bit RSA signature in 1.3 ms and verify it in 0.07 ms. We emulated the overhead of Cauchy-based Reed-Solomon erasure codes [6] used for reconciliation. Servers and clients sent unicast messages. We used the netem utility to place delay and bandwidth constraints on the links between the servers. We added 50 ms delay (emulating a US-wide deployment) to each link and limited the aggregate outgoing bandwidth of each server to 10 Mbps. Clients were evenly distributed among the servers and no delay or bandwidth constraints were set between the client and its server.

Attack Strategies: Our experimental results during attack show the minimum performance that must be achieved in order for a malicious leader to remain in power. Our measurements do not reflect the time required for view changes, during which a new leader is installed. Since a view change takes a finite, and, in practice, relatively small, amount of time, malicious leaders must cause performance degradation without being detected in order to have a prolonged effect on throughput. Therefore, we focus on the attack scenario where a malicious leader stays in power indefinitely while degrading performance.

We use the first attack on BFT described in Section 3. We present results for a very aggressive yet possible time-




Figure 6: Latency (ms) vs. number of clients, 50 ms Diameter

out (300 ms), yielding the most favorable performance for BFT under attack. To attack Prime, (1) the leader adds as much delay as possible (without being suspected) to the protocol, and (2) faulty servers force as much reconciliation as possible. A malicious leader can add approximately two rounds of delay to the global ordering phase (see Figure 2). The malicious servers (servers 1 and 2) force reconciliation by not sending their PO-REQUEST messages to f of the correct servers (servers 6 and 7). Therefore, all updates originating from the faulty servers must be sent to servers 6 and 7 using the reconciliation mechanism (Section 4). Moreover, the malicious servers only acknowledge each other's PO-REQUEST messages, forcing the correct servers to send reconciliation messages to servers 1 and 2 for all messages originating from correct servers. Thus, all messages undergo a reconciliation step, which consumes approximately the same outgoing bandwidth as update dissemination during preordering. This reduces the maximum achievable throughput by approximately half.

Performance Results: Figure 5 shows system throughput in updates per second as a function of the number of clients. The clients send one write update (containing 512 bytes of data), wait for proof that the update has been ordered, and then submit their next update. BFT uses an optimization where clients send updates directly to all of the servers, and the BFT PRE-PREPARE message contains batches of update digests. When both protocols are not under attack, the throughput of BFT increases at a faster rate than the throughput of Prime, because BFT has fewer protocol rounds. BFT's performance plateaus due to bandwidth constraints at slightly less than 850 updates per second with about 250 clients. Prime reaches a similar plateau with about 350 clients.

Throughput results are much different when the two protocols are attacked. With an aggressive timeout of 300 ms, BFT can order less than 30 updates per second. With the default timeout of 5 sec, BFT can only order 2 updates per second (not shown). Prime plateaus at about 400 updates per second due to the bandwidth overhead of reconciliation. The slope of the curve corresponding to Prime under attack

is less steep than when it is not under attack due to the delay added by the malicious leader. We include results with $K_{Lat} = 1$ and $K_{Lat} = 2$. K_{Lat} accounts for variabilities in latency (Section 2). As K_{Lat} increases, a malicious leader can add more delay to the turn-around time without being detected. Prime’s throughput continues to increase until it becomes bandwidth constrained. BFT reaches its maximum throughput when there is one client per server. This throughput limitation, which occurs when only a small amount of the available bandwidth is used, is a consequence of judging the leader conservatively.

Figure 6 shows update latency, measured at the client, as a function of the number of clients. When the protocols are not under attack, BFT has a lower latency than Prime, due to the differences in the number of protocol rounds. The latency of both protocols increases at different points before the plateau due to overhead associated with aggregation. The latency begins to climb steeply when the throughput plateaus due to update queueing at the servers. When under attack, the latency of Prime increases due to the two extra protocol rounds added by the leader. When $K_{Lat} = 2$, the leader can add approximately 100 ms more delay than when $K_{Lat} = 1$. The latency of BFT under attack climbs as soon as more than one client is added to each server, because the leader can order one update per server during a timeout period without being suspected.

7 Related Work

The protocols considered in this paper use the state machine approach [17, 26] to achieve replication, in which replicas execute a totally ordered stream of updates. This paper focused on leader-based Byzantine fault-tolerant SMR protocols. Other approaches are possible, such as using randomization, quorum systems, and hybrid architectures. A thorough analysis of how resilient these systems are to performance failures is an interesting avenue for future research.

Many Byzantine fault-tolerant SMR systems rely on a leader to coordinate the ordering protocol [2, 4, 8, 16, 18, 21, 29]. The consistency of these systems does not rely on synchrony assumptions, while liveness is guaranteed assuming the network meets certain stability properties. To ensure that the stability properties are eventually met in practice, they use exponentially growing timeouts during view changes. This makes these systems vulnerable to the type of performance degradation when under attack described in Section 3.2. In contrast, Prime uses the SUSPECT-LEADER protocol to allow correct servers to collectively decide whether the leader is performing fast enough by adapting to the network conditions once the system stabilizes.

Rampart [24] implements Byzantine atomic multicast over a reliable group multicast protocol. This is similar

to how Prime uses preordering followed by global ordering. Both protocols disseminate updates to $2f + 1$ servers before a coordinator assigns the global order. Drabkin et al. [12] observe the difficulty of setting protocol timeouts in the context of group communication in malicious settings.

Other Byzantine fault-tolerant protocols [5, 7, 22, 23] use randomization to circumvent the FLP impossibility result, guaranteeing termination with probability 1. These protocols incur a high number of communication rounds during normal-case operation (even those that terminate in an expected constant number of rounds). However, they do not rely on a leader to coordinate the ordering protocol, and thus may not suffer the same kinds of performance vulnerabilities when under attack. We believe it is an interesting open question to consider (1) whether their performance in fault-free configurations is sufficiently high, especially in high latency environments and (2) how resilient they are to performance degradation when under attack.

Byzantine quorum systems [1, 10, 19, 20] can also be used for replication. While early work in this area was restricted to a read/write interface, recent work uses quorum systems to provide SMR. The Q/U protocol [1] of Abd-El-Malek et al. requires $5f + 1$ replicas for this purpose and suffers performance degradation when write contention occurs. The HQ protocol [10] showed how to mitigate this cost by reducing the number of replicas to $3f + 1$. Since HQ uses BFT to resolve contention when it arises, it is vulnerable to the same types of performance degradation as BFT.

A different approach to SMR is to use a hybrid architecture in which different parts of the system rely on different fault and/or timing assumptions [9, 27, 28]. The different components are therefore resilient to different types of attacks. We believe leveraging stronger timing assumptions may allow for more aggressive performance monitoring.

The Θ -Model [15] assumes that messages in transit simultaneously experience a bounded ratio of end-to-end delays. PRIME-STABILITY assumes an eventual ratio of delays on each link between correct servers.

8 Conclusions

In this paper we brought to light the vulnerability of current leader-based Byzantine fault-tolerant SMR protocols to performance degradation when under attack. We proposed the BOUNDED-DELAY correctness criterion to complement current liveness criteria by requiring the leader to act timely in order to stay in power. We presented Prime, a new Byzantine fault-tolerant SMR protocol, which meets BOUNDED-DELAY and is a first step towards making intrusion-tolerant replication resilient to performance attacks in malicious environments. Our experimental results show that Prime performs competitively with BFT in fault-free configurations and an order of magnitude better when under attack.

References

- [1] M. Abd-El-Malek, G. Ganger, G. Goodson, M. Reiter, and J. Wylie. Fault-scalable Byzantine fault-tolerant services. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles (SOSP '05)*, pages 59–74, Brighton, UK, 2005.
- [2] Y. Amir, B. Coan, J. Kirsch, and J. Lane. Customizable fault tolerance for wide-area replication. In *Proceedings of the 26th IEEE International Symposium on Reliable Distributed Systems (SRDS '07)*, pages 66–80, Beijing, China, 2007.
- [3] Y. Amir, B. Coan, J. Kirsch, and J. Lane. Byzantine replication under attack. Technical Report CNDS-2008-1, Johns Hopkins University, www.dsn.jhu.edu, March 2008.
- [4] Y. Amir, C. Danilov, J. Kirsch, J. Lane, D. Dolev, C. Nita-Rotaru, J. Olsen, and D. Zage. Scaling Byzantine fault-tolerant replication to wide area networks. In *Proceedings of the 2006 International Conference on Dependable Systems and Networks (DSN'06)*, pages 105–114, Philadelphia, PA, USA, June 2006.
- [5] M. Ben-Or. Another advantage of free choice (extended abstract): Completely asynchronous agreement protocols. In *Proceedings of the 2nd Annual ACM Symposium on Principles of Distributed Computing (PODC '83)*, pages 27–30, 1983.
- [6] J. Blomer, M. Kalfane, M. Karpinski, R. Karp, M. Luby, and D. Zuckerman. An xor-based erasure-resilient coding scheme. Technical Report TR-95-048, International Computer Science Institute, August 1995.
- [7] C. Cachin and J. A. Portiz. Secure intrusion-tolerant replication on the internet. In *Proceedings of the 2002 International Conference on Dependable Systems and Networks (DSN '02)*, pages 167–176, Bethesda, MD, USA, June 2002.
- [8] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation (OSDI '99)*, pages 173–186. USENIX Association, Co-sponsored by IEEE TCOS and ACM SIGOPS, 1999.
- [9] M. Correia, N. F. Neves, and P. Veríssimo. How to tolerate half less one Byzantine nodes in practical distributed systems. In *Proceedings of the 23rd IEEE International Symposium on Reliable Distributed Systems (SRDS'04)*, pages 174–183, Florianopolis, Brazil, 2004.
- [10] J. Cowling, D. Myers, B. Liskov, R. Rodrigues, and L. Shrira. HQ replication: A hybrid quorum protocol for Byzantine fault tolerance. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI '06)*, pages 177–190, Seattle, WA, Nov. 2006.
- [11] F. Cristian. Understanding fault-tolerant distributed systems. *Communications of the ACM*, 34(2):56–78, 1991.
- [12] V. Drabkin, R. Friedman, and A. Kama. Practical Byzantine group communication. In *Proceedings of the 26th IEEE International Conference on Distributed Computing Systems (ICDCS '06)*, page 36, Lisboa, Portugal, 2006.
- [13] C. Dwork, N. Lynch, and L. Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM*, 35(2):288–323, 1988.
- [14] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.
- [15] J.-F. Hermant and J. Widder. Implementing reliable distributed real-time systems with the θ -model. In *Proceedings of the 9th International Conference on Principles of Distributed Systems (OPODIS '05)*, pages 334–350, Pisa, Italy, December 2005.
- [16] R. Kotla, L. Alvisi, M. Dahlin, A. Clement, and E. Wong. Zyzzyva: speculative Byzantine fault tolerance. In *Proceedings of 21st ACM SIGOPS Symposium on Operating Systems Principles (SOSP '07)*, pages 45–58, 2007.
- [17] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978.
- [18] J. Li and D. Mazieres. Beyond one-third faulty replicas in Byzantine fault tolerant systems. In *Proceedings of the 4th USENIX Symposium on Networked Systems Design and Implementation (NSDI '07)*, pages 131–144, 2007.
- [19] D. Malkhi and M. Reiter. Byzantine quorum systems. *Distributed Computing*, 11(4):203–213, 1998.
- [20] D. Malkhi and M. K. Reiter. Secure and scalable replication in Phalanx. In *Proceedings of the 17th IEEE Symposium on Reliable Distributed Systems (SRDS '98)*, pages 51–58, West Lafayette, IN, USA, 1998.
- [21] J.-P. Martin and L. Alvisi. Fast Byzantine consensus. *IEEE Transactions on Dependable and Secure Computing*, 3(3):202–215, 2006.
- [22] H. Moniz, N. F. Neves, M. Correia, and P. Veríssimo. Randomized intrusion-tolerant asynchronous services. In *Proceedings of the 2006 International Conference on Dependable Systems and Networks (DSN'06)*, pages 568–577, 2006.
- [23] M. O. Rabin. Randomized Byzantine generals. In *The 24th Annual IEEE Symposium on Foundations of Computer Science*, pages 403–409, 1983.
- [24] M. K. Reiter. The Rampart Toolkit for building high-integrity services. In *Selected Papers from the International Workshop on Theory and Practice in Distributed Systems*, pages 99–110, London, UK, 1995. Springer-Verlag.
- [25] L. M. R. Sampaio, F. V. Brasileiro, W. Cirne, and J. C. A. Figueiredo. How bad are wrong suspicions? Towards adaptive distributed protocols. In *Proceedings of the 2003 International Conference on Dependable Systems and Networks (DSN '03)*, pages 551–560, San Francisco, CA, USA, 2003.
- [26] F. B. Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys*, 22(4):299–319, 1990.
- [27] M. Serafini and N. Suri. The fail-heterogeneous architectural model. In *Proceedings of the 26th IEEE International Symposium on Reliable Distributed Systems (SRDS '07)*, pages 103–113, Beijing, China, 2007.
- [28] P. Veríssimo, N. Neves, C. Cachin, J. Poritz, D. Powell, Y. Deswarte, R. Stroud, and I. Welch. Intrusion-tolerant middleware: The road to automatic security. *IEEE Security & Privacy*, 4(4):54–62, 2006.
- [29] J. Yin, J.-P. Martin, A. Venkataramani, L. Alvisi, and M. Dahlin. Separating agreement from execution for Byzantine fault-tolerant services. In *Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP '03)*, pages 253–267, Bolton Landing, NY, USA, October 2003.