Performance Comparison of a Rotating Coordinator and a Leader Based Consensus Algorithm

Péter Urbán[1], Naohiro Hayashibara[1],
André Schiper[2], and Takuya Katayama[1]

[1]School of Information Science, Japan Advanced Institute of Science and Technology
[2]Swiss Federal Institute of Technology in Lausanne (EPFL)

March 30, 2004
IS-RR-2004-008

Research Report

JAIST

School of Information Science
Japan Advanced Institute of Science and Technology

ISSN 0918-7553

Performance Comparison of a Rotating Coordinator and a Leader Based Consensus Algorithm

JAIST Research Report IS-RR-2004-008

Péter Urbán*, Naohiro Hayashibara*, André Schiper†, Takuya Katayama*
*School of Information Science, Japan Advanced Institute of Science and Technology
1-1 Asahidai, Tatsunokuchi, Nomi, Ishikawa 923-1292, Japan
Email: {urban,nao-haya,katayama}@jaist.ac.jp
†Faculté Informatique et Communications, École Polytechnique Fédérale de Lausanne (EPFL)
CH-1015 Lausanne, Switzerland
Email: andre.schiper@epfl.ch

Abstract— **Protocols that solve agreement problems are essential building blocks for fault tolerant distributed systems. While many protocols have been published, little has been done to analyze their performance, especially the performance of their fault tolerance mechanisms. In this paper, we compare two well-known asynchronous consensus algorithms. In both algorithms, a leader process tries to impose a decision, and another leader retries if the leader fails doing so. The algorithms elect leaders differently: the Chandra-Toueg algorithm has a rotating leader, whereas processes in the Paxos algorithm elect leaders directly. We investigate the performance implications of this difference.**

In the system under study, processes send atomic broadcasts to each other. Consensus is used to decide the delivery order of messages. We evaluate the steady state latency in (1) runs with neither crashes nor suspicions, (2) runs with crashes and (3) runs with no crashes in which correct processes are wrongly suspected to have crashed, as well as the transient latency after (4) one crash and (5) multiple correlated crashes. The results show that the Paxos algorithm tolerates frequent wrong suspicions (3) and correlated crashes (5) better, while the performance is comparable in all other scenarios.

Keywords: simulation, consensus, atomic broadcast, rotating coordinator, leader, asynchronous, failure detector

I. Introduction

Agreement problems — such as consensus, atomic broadcast or atomic commitment — are essential building blocks for fault tolerant distributed applications, including transactional and time critical applications. These agreement problems have been extensively studied in various system models, and many protocols solving these problems have been published [1], [2], offering different levels of guarantees. However, these protocols have mostly been analyzed from the point of view of their safety and liveness properties, and very little has been done to analyze their *performance*. Also, most papers focus on analyzing failure free runs, thus neglecting the performance aspects of failure handling. In our view, the limited understanding of performance aspects, in both failure free scenarios and scenarios with failure handling, is an obstacle for adopting such protocols in practice. This paper presents a performance study focusing on consensus, a problem related to most other agreement problems [3], in scenarios that involve failure handling.

1) The two algorithms: We present a study comparing the performance of two consensus algorithms: the Chandra-Toueg [4] and Paxos [5], [6] algorithms. We have chosen these algorithms because they are well-known, and because there is an ongoing informal debate in the community about their relative performance. We hope that our comparison will bring some objective arguments to this debate.

The algorithms share a number of characteristics: they are designed for the asynchronous system model with minimal extensions needed to solve consensus, they need that a majority of processes is correct, and they have a similar structure: they execute a sequence of rounds whereby each round has a leader[1] that tries to impose a decision. They differ in how they tolerate (suspected) failures of the leader: processes in the Chandra-Toueg algorithm rotate the leader role among all processes, whereas processes in the Paxos algorithm elect leaders directly in an uncoordinated manner (the two approaches are sometimes called rotating coordinator paradigm and leader based paradigm). In this paper, we investigate the performance implications of this difference.

2) Elements of the performance study: The two consensus algorithms are analyzed in a system in which processes send atomic broadcasts to each other. Since the atomic broadcast algorithm that we use [4] leads to the execution of a sequence of consensus to decide the delivery order of messages, evaluating the performance of atomic broadcast is a good way of evaluating the performance of the underlying consensus algorithm in a realistic usage scenario. In our study, the atomic broadcast algorithm uses either of the two consensus algorithms. We study the system using simulation, which allows us to compare the algorithms in a variety of different environments. We model message exchange by taking into account contention on the network and the hosts, using the

Research supported by the Japan Society for the Promotion of Science and the CSEM Swiss Center for Electronics and Microtechnology, Inc., Neuchâtel.

[1] Ref. [4] uses the term *coordinator*. We stick to *leader* throughout the paper.

metrics described in [7], [8]. We model failure detectors in an abstract way, using the quality of service (QoS) metrics proposed by Chen et al. [9]. We compare the algorithms using the benchmarks proposed in [8], [10] (which are stated in terms of the system under study, i.e., atomic broadcast). Our main performance metric for atomic broadcast is *early latency*, the time that elapses between the sending of a message m and the earliest delivery of m. We use symmetric workloads. We evaluate the steady state latency in (1) runs with neither crashes nor suspicions, (2) runs with crashes and (3) runs with no crashes in which correct processes are wrongly suspected to have crashed, as well as the transient latency after (4) one crash and (5) multiple correlated crashes.

3) The results: Our main finding is that, although the two algorithms have comparable performance in scenarios (1), (2) and (4), *the Paxos algorithm performs significantly better in scenarios 3 and 5*. With multiple correlated crashes, the reason is that the Paxos algorithm elects a correct leader immediately after detecting the crashes. We found the largest difference when wrong failure suspicions were frequent and/or long lasting wrong failure suspicions. The reason is that the Paxos algorithm generates less contention: its leader election mechanism makes sure that only a small subset of all processes start concurrent rounds, whereas the rotating leader scheme in the Chandra-Toueg algorithm results in nearly all processes starting concurrent rounds. Therefore the leader based approach seems more suited to environments in which the failure detection service makes mistakes often.

4) Structure: The rest of the paper is structured as follows. Section II presents related work. Section III defines the system model and the agreement problems used in this paper. We introduce the algorithms in Section IV. Section V describes the benchmarks we used, followed by our simulation model for the network and the failure detectors in Section VI. Our results are presented in Section VII, and the paper concludes with a discussion in Section VIII.

II. Related work

Most of the time, consensus algorithms are evaluated using simple metrics like time complexity (number of communication steps) and message complexity (number of messages). This gives, however, little information on the real performance of those algorithms. A few papers provide a more detailed performance analysis. Ref. [11] compares the impact of different implementations of failure detectors on the Chandra-Toueg consensus algorithm; Ref. [12] and [13] analyze the latency of the same algorithm, concentrating mostly on the effect of wrong failure suspicions; All these papers consider only isolated consensus executions, which are a special case of our workloads, corresponding to a very low setting for the throughput. Other papers [10], [15] consider a consensus algorithm embedded in an atomic broadcast algorithm, but they do not aim at comparing consensus algorithms. Note also that the performance of atomic broadcast algorithms is studied more extensively in the literature than the performance of consensus algorithms (see [8] for a summary).

Most papers on the performance of agreement algorithms only consider failure free executions (our normal-steady faultload), which only gives a partial and incomplete understanding of the behavior of the algorithms. We only note a few interesting exceptions here. The transient effects of a crash are studied in [10], [11], [16], but the faultload in [11], [16] is different from our crash-transient faultload. Ref. [11] assumes that the crash occurs at the worst possible moment during execution, leading to the worst case latency. In contrast to our faultload, this faultload requires a detailed knowledge of the execution, which is only available if one considers very simple workloads (isolated executions of consensus in [11]) in an analytical or simulation model. The other paper [16] measures the latency of the group membership service used by the algorithm to tolerate crash failures.[2] This way of considering the transient effects of a crash is less general compared to our faultload, as it is stated in terms of an implementation detail of the algorithm under study.

The assumptions and/or the algorithms used in all the studies listed are too different to allow a meaningful comparison of the results with those in this paper. Our previous work [14] would be an exception: it compares the same algorithms using measurements rather than simulation, and with fewer faultloads. However, bugs discovered and fixed since its publication invalidate the results presented there.

III. Definitions

A. System model

We consider a widely accepted system model. It consists of n processes $p_1, \ldots, p_n$ that communicate only by message passing. The system is asynchronous, i.e., we make no assumptions on its timing behavior: there are no bounds on the message transmission delays and the relative processing speeds of processes. The network is quasi-reliable: it does not lose, alter nor duplicate messages (messages whose sender or recipient crashes might be lost). In practice, this is easily achieved by retransmitting lost messages. We consider that processes only fail by crashing. Crashed processes do not send any further messages. Process crashes are rare, processes fail independently, and process recovery is slow: both the time between crashes and time to repair are much greater than the latency of the algorithms investigated.

The consensus algorithms used in this paper use *oracles* to tolerate process crashes: the Chandra-Toueg algorithm (CT) uses *failure detector oracles* and the Paxos algorithm (Paxos) uses *leader oracles*. A failure detector oracle outputs a list of processes it suspects to have crashed. It might make mistakes: it might suspect correct processes and it might not suspect crashed processes immediately. A leader oracle outputs a single leader process that it trusts to be alive. All leader oracles in the system strive to output the same leader process. This oracle might make mistakes as well: it might elect crashed processes as leader, and different oracles might elect different leaders. To make sure that the consensus algorithms terminate,

[2] Certain kinds of Byzantine failures are also injected.

we need some assumptions on the behavior of the oracles: $\Diamond\mathcal{S}$ for CT [4] and Ω for Paxos [17]. These assumptions are rather weak: they can usually be fulfilled in real systems by tuning implementation parameters of the oracles [15], [18]. Also, they are equivalent: one can solve the same set of problems when using the asynchronous model with oracles fulfilling either of $\Diamond\mathcal{S}$ and Ω [17].

B. Agreement problems

We next give informal definitions of the agreement problems needed for understanding this paper; see [4], [19] for more formal definitions.

In the consensus problem, each process proposes an initial value. Uniform consensus (considered here) ensures that all processes decide the same value, which is one (any one) of the proposals.

Atomic broadcast is defined in terms of two primitives called *A-broadcast*(m) and *A-deliver*(m), where m is some message. Uniform atomic broadcast (considered here) guarantees that (1) if a message is A-broadcast by a process, then all correct processes eventually A-deliver it, and (2) all processes A-deliver messages in the same order.

The algorithms in this study use (non-uniform) reliable broadcast, which guarantees that if a message is broadcast by a correct process, then all correct processes eventually deliver it (even if the sender crashes).

IV. Algorithms

This section sketches the two consensus algorithms, concentrating on their common points and their differences. We then introduce the atomic broadcast algorithm built on top of consensus.

A. The consensus algorithms

For solving consensus, we use the Chandra-Toueg $\Diamond\mathcal{S}$ algorithm [4] and the single-decree Synod algorithm from the Paxos paper [5], [6]. Henceforth, we shall refer to the algorithms as *CT algorithm* and *Paxos algorithm*, respectively. We also use these names to refer to the atomic broadcast algorithm used with the corresponding consensus algorithm if no confusion arises from doing so.

1) Common points: The algorithms share a lot of assumptions and characteristics, which makes them ideal candidates for a performance comparison. In particular, both algorithms are designed for the asynchronous model with equally strong oracles: $\Diamond\mathcal{S}$ failure detectors (CT algorithm; see Section III-A) and Ω leader oracles (Paxos algorithm). Both tolerate $f < n/2$ crash failures. In both algorithms, processes execute a sequence of asynchronous rounds (i.e., not all processes necessarily execute the same round at a given time t). Each round has a *leader* (called *coordinator* in [4]), whose role is to try to impose a decision value on all processes. If it succeeds, the consensus algorithm terminates; if it fails, some additional rounds are executed with possibly a different leader. Moreover, leaders execute a very similar protocol in each round,[3] discussed in detail in Section IV-A.3.

2) Electing a leader: The main difference between the algorithms is how the leaders are chosen. A new leader is necessary whenever the current round is not successful. A round may not be successful if one or more processes want a different leader, usually because they suspect the current leader to have crashed.

The CT algorithm is based on the rotating coordinator paradigm. Whenever the current leader is suspected, the leader is chosen to be the next process, in a round-robin fashion. In other words, each process executes a sequence of rounds $1, 2, \ldots$, and there is *a priori agreement* on the identity of the leader: process p_i is leader for rounds $kn + i$.

There is no such a priori agreement in the Paxos algorithm. A process p_i considers itself leader (and starts a new round) when its leader oracle outputs p_i. Other processes only start participating in this round when they receive a message from the leader. Leaders always choose unique increasing round numbers: process p_i is leader for rounds $kn + i$, just like in the CT algorithm. However, unlike in the CT algorithm, a given process hardly ever executes all of the rounds $1, 2, \ldots$: there are usually gaps in the sequence of rounds.

3) Execution of a round: We now sketch the execution of one round in each of the two algorithms, illustrated in Fig. 1. Further details of the execution are not necessary for understanding the rest of the paper.

a) Read phase: Throughout the execution, processes maintain their current *estimate* of the decision value. Both algorithms start the round with a *read phase* whose purpose is to update the leader's estimate with a recent estimate. In the Paxos algorithm, the leader sends a *read* message to all processes, and all processes reply with their estimate (*estimate* messages). In the CT algorithm, the read message is not necessary, as all processes execute every round. In each of the two algorithms, the leader only waits for an estimate from a majority of all processes, and then updates its own estimate.

b) Write phase: In this phase, the leader sends its estimate to all, proposing its acceptance (*proposal* messages). A process accepts this estimate if it has not seen messages from a later round (in the case of the Paxos algorithm) or if it does not suspect the leader (CT algorithm).

When a process accepts a proposal, it updates its own estimate and sends back an *ack* message; otherwise, it sends back a *nack* message (not shown in Fig. 1). In the case of the CT algorithm, the nack message is sent *before* receiving the proposal.

The leader waits for messages from a majority of all processes, and decides if it has received a majority of ack messages. In this case, it also sends a *decision* message to all using reliable broadcast. Upon receiving this message, the other processes decide as well. If the leader receives one nack message before deciding (this is not shown in Fig. 1) it finishes executing the round without deciding.

[3] This is why we chose the CT algorithm over other algorithms written for the same system model (e.g., [20]).

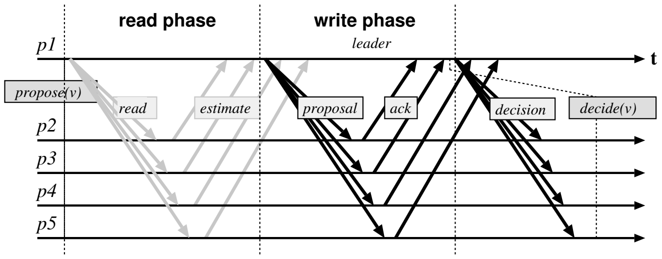

Fig. 1. Example of a round in the CT and Paxos algorithms (CT does not send the *read* message)

B. Optimizations to the consensus algorithms

The consensus algorithms implemented contain several optimizations with respect to the published versions [4]–[6]. The goal of the optimizations is to reduce the number of messages in the most common scenario: when no process is suspected (CT algorithm) or when the leader is the same process (p_1) throughout the execution (Paxos algorithm).

- The read phase is not necessary in the first round, in either of the two algorithms. This is why its messages are gray in Fig. 1.
- In the original CT algorithm, the non-leader processes start the next round immediately after sending the *ack* message. This generates *estimate* messages which are not needed in the most common scenario. These messages degrade performance. To prevent this, non-leader processes wait for an *abort* message before starting the new round.[4] The *abort* message is sent by the leader if it receives *nack* messages.
- In the write phase, the leader stops the current round after receiving the first *nack* message, because it is known at this point already that the round has failed. The original algorithms always wait for (*ack* and *nack*) messages from a majority of processes.
- In both algorithms, the decision message must be sent using reliable broadcast (see Section III-B). We use an efficient algorithm inspired by [21] that requires only one broadcast message if the sender is not suspected.
- The CT algorithm always starts with the same leader p_1. If p_1 crashes, this affects steady-state performance negatively. We fix this problem by having the consensus decide on the first leader of the next consensus (beside the order of messages) [22]. Processes propose the first process that their failure detector trusts as first leader. This choice makes sure that, eventually, crashed processes do not ever become first leaders.

C. The Chandra-Toueg atomic broadcast algorithm

In the Chandra-Toueg atomic broadcast algorithm [4], a process executes A-broadcast by sending a message to all processes.[5] When a process receives such a message, it buffers it until the delivery order is decided. The delivery order is decided by a sequence of consensus numbered 1, 2, etc. The value proposed initially and the decision value of each consensus are *sets of message identifiers*. Let $\mathrm{msg}(k)$ be the set of message IDs decided by consensus $\#k$. The messages denoted by $\mathrm{msg}(k)$ are A-delivered before the messages denoted by $\mathrm{msg}(k+1)$, and the messages denoted by $\mathrm{msg}(k)$ are A-delivered according to a deterministic function, e.g., according to an order relation defined on their IDs.

The algorithm inherits the system model and any fault tolerance guarantees from the underlying consensus algorithm. We use this atomic broadcast algorithm with both the CT and Paxos consensus algorithms.

The performance of the algorithms can be improved by packing messages from subsequent consensus executions into one message. For the sake of simplicity, we did not perform such optimizations [23]–[25]. This decision affects the two algorithms in the same way, hence we introduce no bias in the performance study.

V. Benchmarks

This section describes our benchmarks, consisting of performance metrics, workloads and faultloads. In order to get meaningful results, we state the benchmarks in terms of the system under study (processes sending atomic broadcasts) rather than in terms of the component under study (consensus). Previous versions of the benchmarks are published in [8], [10].

A. Performance metrics and workloads

Our main performance metric is the *early latency* of atomic broadcast. Early latency L is defined for a single atomic broadcast as follows. Let *A-broadcast(m)* occur at time t_0, and *A-deliver(m)* on p_i at time t_i, for each $i = 1, \dots, n$. Then latency is defined as the time that elapses until the first A-delivery of m, i.e., $L \overset{\text{def}}{=} (\min_{i=1,\dots,n} t_i) - t_0$. In our study, we compute the mean for L over a lot of messages and several executions.

This performance metric makes sense in practice. Consider a service replicated for fault tolerance using active replication [26]. Clients of this service send their requests to the server replicas using Atomic Broadcast. Once a request is delivered, the server replica processes the client request, and sends back a reply. The client waits for the first reply, and discards the other ones (identical to the first one). If we assume that the time to service a request is the same on all replicas, and the time to send the response from a server to the client is the same for all servers, then the first response received by the client is the response sent by the server to which the request was delivered first. Thus there is a direct link between the response time of the replicated server and the latency L.

Beside the early latency, we also compute the *late latency*, the time that elapses until the last A-delivery of a message m: $L_{late} \overset{\text{def}}{=} (\max_{i=1,\dots,n} t_i) - t_0$.

Latency is always measured under a certain workload. We chose simple workloads: (1) all destination processes send atomic broadcast messages at the same constant rate, and

[4] The non-leader processes also start a new round if they start suspecting the leader.

[5] This message is sent using reliable broadcast. We use the efficient algorithm mentioned Section IV-B.

(2) the A-broadcast events come from a Poisson stochastic process. We call the overall rate of atomic broadcast messages *throughput*, denoted by T. In general, we determine how the latency L depends on the throughput T.

The system can only reach a steady state if the throughput is under some maximal throughput T_{max}. Beyond this throughput, some processes are left behind. We detect if the system reaches steady state by observing if the late latency stabilizes over time.

B. Faultloads

The faultload is the part of the workload that describes failure-related events that occur during an experiment [27]. We concentrate on (1) crash failures of processes, and (2) the behavior of unreliable failure detectors. We evaluate the performance of the algorithms with four different faultloads. We now describe each of them in detail, mentioning which parameters influence latency with each faultload.

1) Normal-steady faultload: With this faultload, we have neither crashes nor wrong suspicions in the experiment. We measure latency after the system reaches its steady state (a sufficiently long time after startup). Parameters that influence latency under this faultload are the algorithm (A), the number of processes (n) and the throughput (T).

2) Crash-steady faultload: One or more crashes occur before the experiment. We measure latency after the system reaches its steady state: a sufficiently long time after startup and any crashes. Beside A, n and T, an additional parameter is the set of crashed processes. In the steady state of the system, all failure detectors in the system permanently suspect all crashed processes at this point, and all leader oracles have elected the same correct process. No wrong suspicions occur, and the leader no longer changes.

3) Crash-transient faultload: With this faultload, we inject one or more crashes at some point in time after the system reached a steady state. Multiple crashes represent correlated failures. After the crashes, we can expect a halt or a significant slowdown of the system for a short period. We would like to capture how the latency changes in atomic broadcasts directly affected by the crashes. Our faultload definition represents the simplest possible choice: we determine the latency of an atomic broadcast sent at the moment of the crashes (by a process that does not crash). Of course, the latency of this atomic broadcast may depend on the choice for the sender and the crashing processes. In order to reduce the number of parameters, we consider the worst case, i.e., the case that increases latency the most.

The precise definition for the faultload is the following. Consider that a set of c processes $\mathcal{C}$ crashes at time t (no other crashes nor wrong suspicions occur). Let process p ($p \notin \mathcal{C}$) execute *A-broadcast*(m) at t. Let $L(p, \mathcal{C})$ be the mean latency of m, averaged over a lot of executions. Then $L_{crash} \overset{\text{def}}{=} \max_{p,\mathcal{C}} L(p, \mathcal{C})$, i.e., we choose the sender and the crashing processes such that latency increases the most.

Beside A, n, T and c, an additional parameter describes how fast failure detectors and leader oracles detect the crashes. This parameter is discussed in Section VI-B.

4) Suspicion-steady faultload: No crashes occur, but failure detectors generate wrong suspicions, and leader oracles change their mind about the leader. This causes the algorithms to take extra steps and thus increase latency. Beside A, n and T, additional parameters include how often wrong suspicions occur and how long they last. These parameters are discussed in Section VI-B.

VI. Simulation models

Our approach to performance evaluation is simulation, which allowed for more general results as would have been feasible to obtain with measurements in a real system (we can use a parameter in our network model to simulate a variety of different environments). We used the Neko prototyping and simulation framework [28] to conduct our experiments. We used the same models for our previous work [8], [10].

A. Modeling the execution environment

We now describe how we modeled the transmission of messages. We use a model inspired from simple models of Ethernet networks [29], and validated in [8]. The key point in the model is that it accounts for *resource contention*. This point is important as resource contention is often a limiting factor for the performance of distributed algorithms. Both a host and the network itself can be a bottleneck. These two kinds of resources appear in the model (see Fig. 2): the network resource (shared among all processes) represents the transmission medium, and the CPU resources (one per process) represent the processing performed by the network controllers and the layers of the networking stack, during the emission and the reception of a message (the cost of running the algorithm is negligible). A message m transmitted for process p_i to process p_j uses the resources (1) CPU_i, (2) network, and (3) CPU_j, in this order. Message m is put in a waiting queue before each stage if the corresponding resource is busy. The time spent on the network resource is one time unit. The time spent on each CPU resource is λ time units; the underlying assumption is that sending and receiving a message has a roughly equal cost.

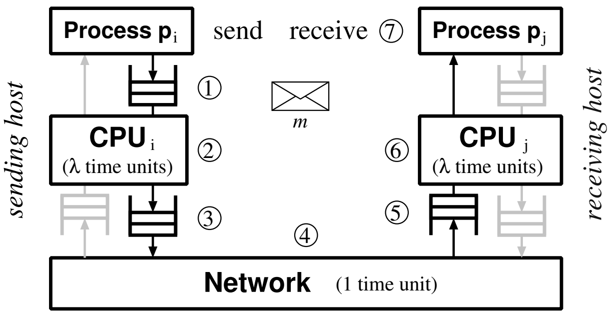

Fig. 2. Transmission of a message in our network model.

The λ parameter ($0 \leq \lambda$) shows the relative speed of processing a message on a host compared to transmitting it over the network. Different values model different networking environments. We conducted experiments with a variety of settings for λ.

We model network-level multicasts: a message sent to several destinations is only processed once on the sending CPU resource and on the network resource.

Crashes are modeled as follows. If a process p_i crashes at time t, no messages can pass between p_i and CPU_i after t; however, the messages on CPU_i and the content of the attached queues are still sent, even after time t. In real systems, this corresponds to a (software) crash of the application process (operating system process), rather than a (hardware) crash of the host or a kernel panic. We chose to model software crashes because they are more frequent in most systems [30].

B. Modeling failure detectors

One approach to examine the behavior of a failure detector is implementing it and using the implementation in the experiments. However, this approach would restrict the generality of our performance study: another choice for the algorithm would likely give different results. Also, it is not justified to model the failure detector in so much detail, as other components of the system, like the execution environment, are modeled much more coarsely. We built a more abstract model instead, using the notion of quality of service (QoS) of failure detectors introduced in [9]. The authors consider the failure detector at a process q that monitors another process p, and identify the following three primary QoS metrics:

- *Detection time T_D:* The time that elapses from p's crash to the time when q starts suspecting p permanently. The definition is illustrated in Fig. 3.
- *Mistake recurrence time T_{MR}:* The time between two consecutive mistakes (q wrongly suspecting p), given that p did not crash; see Fig. 4.
- *Mistake duration T_M:* The time it takes a failure detector component to correct a mistake, i.e., to trust p again (given that p did not crash); see Fig. 4.

Not all of these metrics are equally important in each of our faultloads (see Section V-B). In the *normal-steady* faultload, the metrics are not relevant. The same holds in the *crash-steady* faultload, because we observe the system a sufficiently long time after all crashes, long enough to have all failure detectors to suspect the crashed processes permanently. In the *suspicion-steady* faultload no crash occurs, hence the latency of atomic broadcast only depends on T_{MR} and T_M (shown in Fig. 4). In the *crash-transient* faultload no wrong suspicions occur, hence T_D is the relevant metric (shown in Fig. 3).

In [9], the QoS metrics are random variables, defined on a pair of processes. In our system, where n processes monitor each other, we have thus $n(n-1)$ failure detectors in the sense of [9], each characterized with three random variables (T_D, T_{MR}, T_M). In order to have an executable model for the failure detectors, we have to define (1) how these random variables depend on each other, and (2) how the distribution of each random variable can be characterized. To keep our model simple, we assume that all failure detector modules are independent and the tuples of their random variables are identically distributed. Moreover, note that we do not need to model how T_{MR} and T_M depend on T_D, as the two former are only relevant in the suspicion-steady faultload, whereas T_D is only relevant in the crash-transient faultload. As for the distributions of the metrics, we took the simplest possible choices: T_D is a constant, and both T_{MR} and T_M are exponentially distributed with (different) constant parameters. This choice only represents a starting point, as we are not aware of any previous work we could build on (apart from [9] that makes similar assumptions). We will refine our models as we gain more experience.

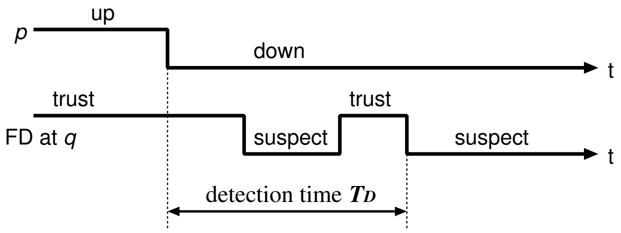

Fig. 3. Quality of service metric expressing the speed of failure detection. Process q monitors process p.

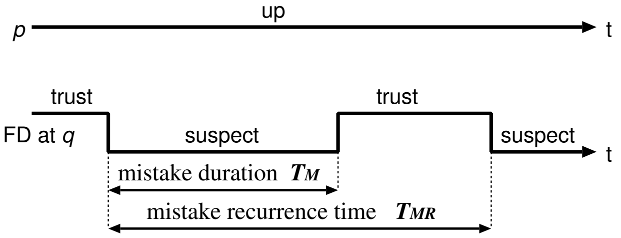

Fig. 4. Quality of service metrics describing wrong suspicions made by failure detectors. Process q monitors process p.

Finally, note that this abstract model for failure detectors neglects that failure detectors and their messages put a load on system components. This simplification is justified in a variety of systems, in which a rather good QoS can be achieved with failure detectors that send messages infrequently.

C. Modeling leader oracles

Our leader oracles for the Paxos algorithm rely on failure detectors: at any point in time, the leader is the process with the smallest index of all processes trusted by the failure detector. We implemented leader oracles with failure detectors because a leader oracle must detect the failures of the desired leader, and it seems logical to use a service dedicated to failure detection to implement this aspect of the leader oracle.[6]

The failure detectors underlying the leader oracles are modeled with their quality of service parameters as described in the previous section.

[6] The leader oracle has other potential uses, e.g., it can be used to implement load balancing among all correct processes (see Section VIII). We intend to investigate this aspect in the future.

VII. Results

We now present our results for all four faultloads and a variety of network models. We obtained results at a variety of representative settings for λ: 0.1, 1 and 10. The settings $\lambda = 0.1$ and 10 correspond to systems where communication generates contention mostly on the network (at $\lambda = 0.1$) and the hosts (at $\lambda = 10$), respectively, while $\lambda = 1$ is an intermediate setting. For example, in current LANs, the time spent on the hosts is much higher than the time spent on the wire, and thus $\lambda = 10$ is probably the setting that corresponds best to such an environment.

Most graphs show the early latency vs. the throughput. Graphs showing the late latency are presented in the appendix only. Values of the late latency are slightly higher, but all other characteristics of the corresponding graphs are very similar. The reason is that if one process reaches a decision in either of the consensus algorithms, all other processes will soon follow, thanks to the decision message (see Section IV-A.3). The maximal throughput is approximately the highest throughput value, that is, the x coordinate of the rightmost point, in all graphs showing the steady-state latency; beyond this throughput, the late latency did not stabilize (see Section V-A). We set the time unit of the network simulation model to 1 ms, to make sure that the reader is not distracted by an unfamiliar presentation of time/frequency values (one that refers to time units). Any other value could have been used. The 95% confidence interval is shown for each point in the graphs.

The two algorithms were always run with an odd number of processes. The reason is that the same number of crash failures k ($k = 1, 2, \ldots$) is tolerated if the algorithms are run with $2k + 1$ and $2k + 2$ processes; thus adding a process to a system with an odd number of processes does not increase the resiliency of the system.

A. Normal-steady and crash-steady faultloads Figures 5 and 6, Appendixes I and II)

With these faultloads, the two algorithms have the same performance. Each curve thus shows the latency of *both* algorithms. For the sake of readability, we only present a subset of the results in Fig. 5 (normal-steady faultload) and Fig. 6 (crash-steady faultload). The full set of results is presented in Appendixes I and II. The latency increases with the throughput and with the number of processes. Somewhat surprisingly, the latency *decreases* with the number of crashes. The reason is that the crashed processes no longer load the network with messages.

The fact that the two algorithms have the same performance is not surprising. Their only important difference is the way of electing a new leader, and no new leader is elected with these faultloads (such that this influences the steady-state performance). In fact, we have deliberately chosen similar algorithms for this study, so that we can concentrate on the performance differences observed with the other faultloads.

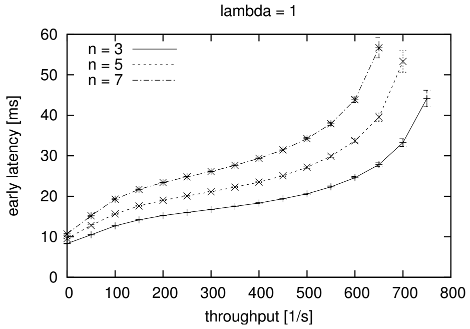

Fig. 5. Latency vs. throughput with the normal-steady faultload, for both Paxos and CT.

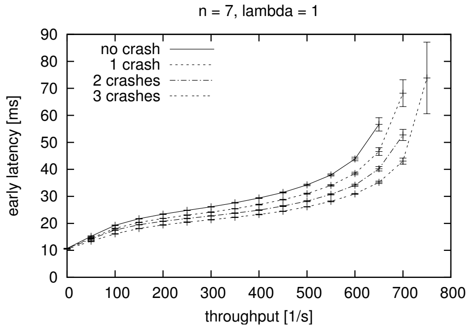

Fig. 6. Latency vs. throughput with the crash-steady faultload, for both Paxos and CT.

B. Crash-transient faultload (Figures 7 and 8, Appendix III)

With this faultload and c crashes, we only present the latency after crashing the first c processes ($p_1, \ldots, p_c$), as this is the case resulting in the highest transient latency (and the most interesting comparison). The crash of any additional processes affects the two algorithms in the same way (slightly decreased latency; cf. Fig. 6).

We set the failure detection timeout T_D to 100 ms at $\lambda = 0.1$ or 1, and to 1000 ms at $\lambda = 10$. This choice models a reasonable trade-off for the failure detector. On the one hand, the detection time T_D is low enough (comparable to the latency overhead) to make sure that the failure detector does not degrade performance catastrophically when a crash occurs. On the other hand, the detection time is high enough (it is a high multiple of the roundtrip time at low loads: $2+4\lambda$) to avoid that failure detectors suspect correct processes.[7]

Both figures show the *latency overhead*, i.e., the latency minus the detection time T_D, rather than the latency. Graphs showing the latency overhead are more illustrative; note that the latency is always greater than the detection time T_D with this faultload, as no atomic broadcast can finish until the crash of the first leader is detected. We start by discussing the results for the case of one crash, shown in Fig. 7. The

[7] As we use an abstract model for the failure detectors for the sake of generality, this does not appear directly in our simulations. The argumentation is about a hypothetical implementation. Given that this implementation can afford spending a high multiple of the roundtrip time before generating a suspicion, wrong suspicions will be rare.

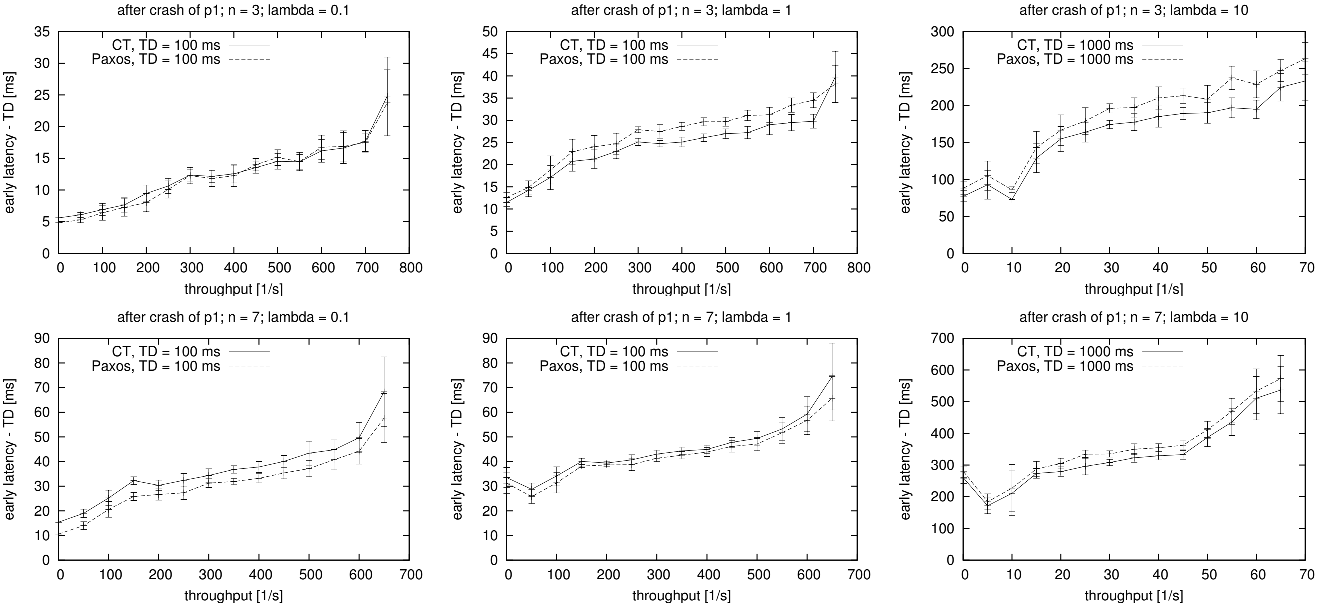

Fig. 7. Latency overhead vs. throughput with the crash-transient faultload. One process crashes.

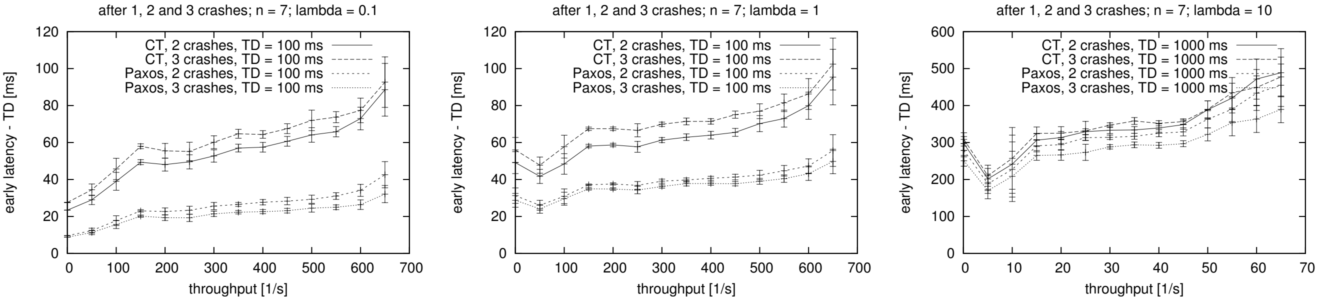

Fig. 8. Latency overhead vs. throughput with the crash-transient faultload. Multiple processes crash.

latency overhead of both algorithms is shown at $n = 3$ (top) and $n = 7$ (bottom) and a variety of values for λ (0.1, 1 and 10 from left to right).

The results show that (1) both algorithms perform rather well (the latency overhead of both algorithms is only a few times higher than the latency with the normal-steady faultload; see Fig. 5) and that (2) the algorithms perform roughly the same. The CT algorithm performs slightly better at $n = 3, \lambda \geq 1$ and $n = 7, \lambda = 10$, i.e., with a small number of processes and a high λ meaning a relatively fast network. The Paxos algorithm performs slightly better at $n = 7, \lambda \leq 1$, i.e., with a lot of processes and a small λ meaning a relatively slow network.

The differences can be explained by differences in the execution of the algorithms once the crash of the first leader is detected. In the CT algorithm, all processes send a *nack* message to the first leader. In the Paxos algorithm, the new leader sends a *read* message. The rest of the execution (from the *estimate* message of the second round) is the same. The CT algorithm thus uses fewer communication steps, but generates more contention on the network; moreover, the increase in network contention is proportional to the number of processes. This explains why the CT algorithm is favored by a fast network and a small number of processes.

The results in the case of multiple crashes are shown in Fig. 8. The latency overhead of both algorithms is shown at $n = 7$, for 2 and 3 crashes (the algorithms do not tolerate more than 3 crashes) and a variety of values for λ (0.1, 1 and 10 from left to right).

The results are different from those obtained with one crash only: the Paxos algorithm always outperforms the CT algorithm. The reason is that the CT algorithm takes more rounds: it rotates over all crashed processes first, whereas the Paxos algorithm elects a correct leader after the first round.

The fact that the CT algorithm rotates over the crashed processes also explains why its latency increases with the number of crashes. The latency of the Paxos algorithm, however, decreases with the number of crashes. The reason is that fewer correct processes load the system with messages to a smaller extent (cf. Fig. 6).

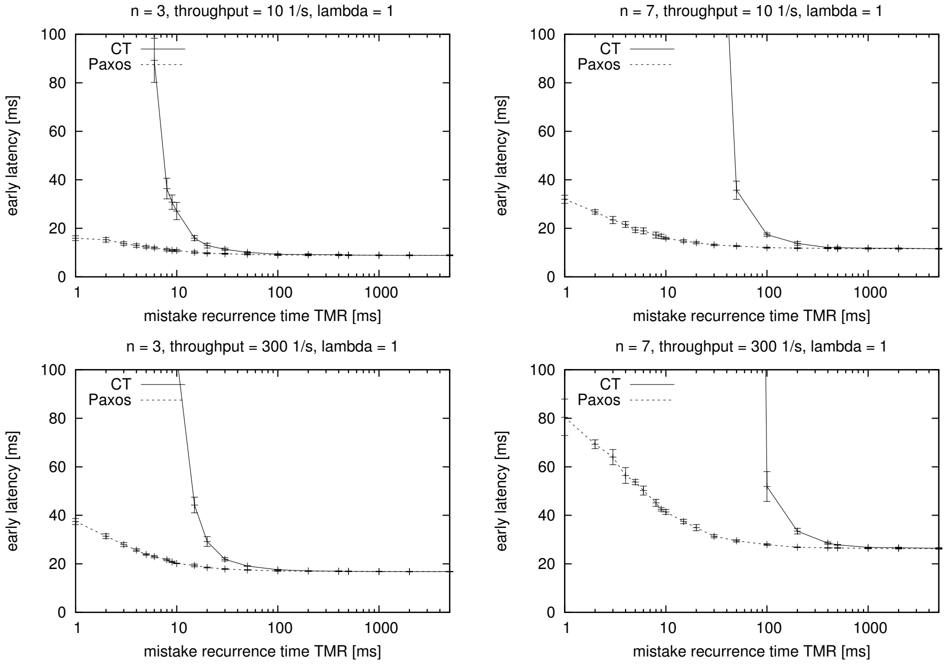

Fig. 9. Latency vs. T_{MR} with the suspicion-steady faultload, with $T_M = 0$ ($\lambda = 1$).

C. Suspicion-steady faultload (Figures 9 and 10, Appendix IV)

The occurrence of wrong suspicions are quantified with the T_{MR} and T_M QoS metrics of the failure detectors. As this faultload does not involve crashes, we expect that the mistake duration T_M is short. In our first set of results (Fig. 9 for $\lambda = 1$; the results for $\lambda = 0.1$ and 10 are similar and are omitted here for better readability; see Appendix IV for the full set of results) we hence set T_M to 0, and latency is shown as a function of T_{MR}. In each figure, we have four graphs: the left column shows results with 3 processes, the right column those with 7; the top row shows results at a low load (10 s^{-1}; 1 s^{-1} if $\lambda = 10$) and the bottom row at a moderate load (300 s^{-1}; 30 s^{-1} if $\lambda = 10$); the algorithms can take a throughput of about 700 s^{-1} (70 s^{-1} if $\lambda = 10$) in the absence of suspicions (i.e., with the normal-steady faultload; see Fig. 5 and Appendix I).

The results show that the CT algorithm is much more sensitive to wrong suspicions if these occur frequently. We illustrate this on Fig. 9: at $n = 3$ and $T = 10$ s^{-1}, that is, the settings at which the CT algorithm tolerates wrong suspicions most, the CT algorithm only works if $T_{MR} \geq 5$ ms, whereas the FD algorithm still works at the smallest T_{MR} value considered (1 ms); the latency of the two algorithms is only equal at $T_{MR} \geq 100$ ms. The CT algorithm breaks down at higher values of T_{MR} for all other settings, whereas the Paxos algorithm continues to work even with 1 ms.

The results can be explained by the difference in the mechanisms that the algorithms use to elect the next leader. The CT algorithm always chooses the next process (in a round-robin manner) as the next leader. Moreover, suspicions are likely to abort the current round. Therefore, if wrong suspicions occur frequently, a lot of rounds are needed to finish a consensus execution, and all processes become leaders, executing rounds that overlap. In contrast, the Paxos algorithm is run with a leader oracle that elects the process with the smallest index among all suspected processes. If suspicions are short ($T_M = 0$), the leader oracle will only ever elect p_1 and p_2 as leader. Only these two processes start overlapping rounds. Moreover, suspicions, even if they lead to a change in the output of the leader oracle, do not abort the current round directly; their only effect is to start other rounds in parallel that might conflict with the current round. Because of these differences, the CT algorithm generates much more contention on the hosts and the network: it is likely that n processes run rounds in parallel, whereas the Paxos algorithm only has two processes that run rounds in parallel. The increased contention of CT is the reason why the Paxos algorithm performs better with this faultload.

In the second set of results (Fig. 10 for $\lambda = 1$; the results for $\lambda = 0.1$ and 10 are similar and are omitted here for better readability; see Appendix IV for the full set of results) T_{MR} is fixed and T_M is on the x axis. We chose T_{MR} such that the latency of the two algorithms is close to equal at $T_M = 0$. For example, with $\lambda = 1$ (Fig. 10), (i) $T_{MR} = 100$ ms for $n = 3$ and (ii) $T_{MR} = 1000$ ms for $n = 7$.

The results show that the CT algorithm is sensitive to the mistake duration T_M as well, not just the mistake recurrence time T_{MR}. Once again, the difference can be attributed to the fact that the Paxos algorithm generates less contention: its leader oracle usually outputs only a small subset of all processes, hence only a few processes start rounds concurrently, whereas all processes are likely to do so in the CT algorithm.

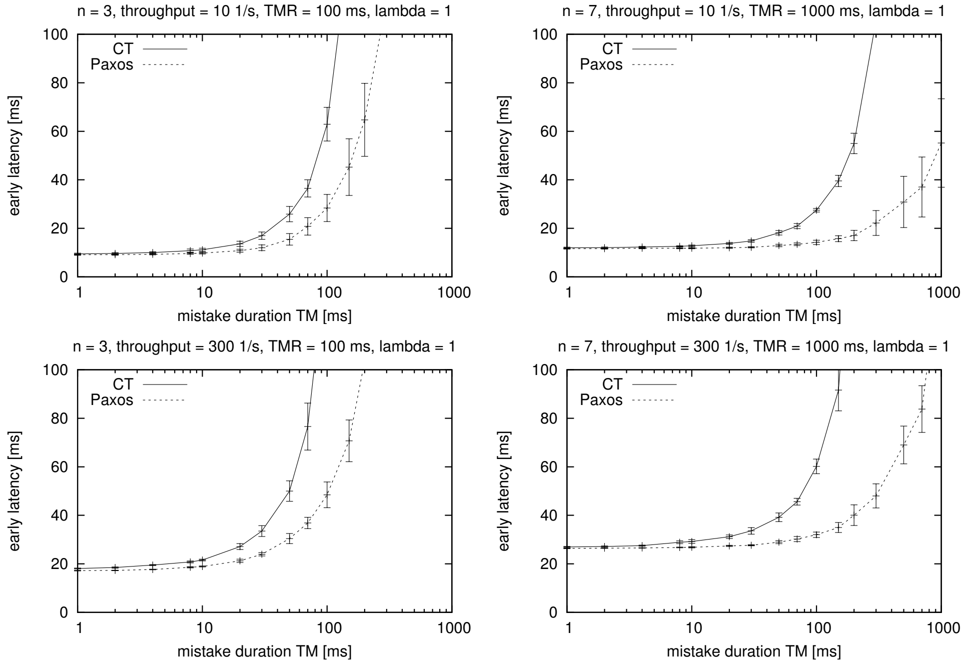

Fig. 10. Latency vs. T_M with the suspicion-steady faultload, with T_{MR} fixed ($\lambda = 1$).

VIII. Discussion

We have compared the performance of the Chandra-Toueg and Paxos consensus algorithms. These algorithms are designed for the same system model: an asynchronous system (with a minimal extension to allow us to solve the consensus problem) and $f < n/2$ process crashes (the highest f that our system model allows). Also, the algorithms have a similar structure: they execute a sequence of rounds whereby each round has a leader that tries to impose a decision. They differ in how they tolerate (suspected) failures of the leader: processes in the Chandra-Toueg algorithm rotate the leader role among all processes, whereas processes in the Paxos algorithm elect leaders directly in an uncoordinated manner.

We found that the two algorithms have the same steady-state performance if neither crashes nor wrong suspicions occur, or if crashes occur but wrong suspicions do not. As for the transient performance after one crash, the performance differences are small, and the relative performance depends on the relative speed of the network and the hosts, as well as on the number of processes. The Paxos algorithm has better transient performance after multiple correlated crashes, because its leader oracle elects a correct leader immediately after detecting the crashes. We found the largest difference in scenarios with frequent or long lasting wrong failure suspicions. In such scenarios, the Paxos algorithm performs better. The reason is that it generates less contention: its leader oracle makes sure that only a small subset of all processes start concurrent rounds, whereas the rotating leader scheme in the Chandra-Toueg algorithm results in nearly all processes starting concurrent rounds. Therefore the leader based approach seems more suited to environments in which the failure detection service makes mistakes often.

We have chosen consensus algorithms with a centralized communication scheme, with one process coordinating the others. In the future, we would like to investigate algorithms with a decentralized communication scheme (e.g., [20]) as well. We would also like to investigate how results change in a load balanced configuration, e.g., in a configuration in which the first leader of subsequent consensus executions rotates among all processes that are alive. The coordinated fashion of electing the next leader in the Chandra-Toueg algorithm might provide performance benefits in such a configuration.

Acknowledgments

We would like to thank Pierre Metrailler for his help in implementing the algorithms and performing the simulations.

References

[1] M. Barborak, M. Malek, and A. Dahbura, “The consensus problem in distributed computing,” *ACM Computing Surveys*, vol. 25, no. 2, pp. 171–220, June 1993.

[2] X. Défago, A. Schiper, and P. Urbán, “Total order broadcast and multicast algorithms: Taxonomy and survey,” Japan Advanced Institute of Science and Technology, Ishikawa, Japan, Research Report IS-RR-2003-009, Sept. 2003.

[3] R. Guerraoui and A. Schiper, “The generic consensus service,” *IEEE Trans. on Software Engineering*, vol. 27, no. 1, pp. 29–41, Jan. 2001. [Online]. Available: http://lsrwww.epfl.ch/Publications/ById/282.html

[4] T. D. Chandra and S. Toueg, “Unreliable failure detectors for reliable distributed systems,” *Journal of ACM*, vol. 43, no. 2, pp. 225–267, 1996.

[5] R. D. Prisco, B. Lampson, and N. Lynch, “Revisiting the PAXOS algorithm,” *Theoretical Computer Science*, vol. 243, no. 1–2, pp. 35–91, July 2000. [Online]. Available: http://www.elsevier.nl/gej-ng/10/41/16/177/21/22/abstract.html; http://www.elsevier.nl/gej-ng/10/41/16/177/21/22/article.pdf

[6] L. Lamport, “The part-time parliament,” *ACM Trans. on Computer Systems*, vol. 16, no. 2, pp. 133–169, 1998.

[7] P. Urbán, X. Défago, and A. Schiper, “Contention-aware metrics for distributed algorithms: Comparison of atomic broadcast algorithms,” in *Proc. 9th IEEE Int’l Conf. on Computer Communications and Networks (IC3N 2000)*, Oct. 2000, pp. 582–589. [Online]. Available: http://lsrwww.epfl.ch/Publications/ById/244.html

[8] P. Urbán, “Evaluating the performance of distributed agreement algorithms: Tools, methodology and case studies,” Ph.D. dissertation, École Polytechnique Fédérale de Lausanne, Switzerland, Aug. 2003, number 2824. [Online]. Available: http://lsrwww.epfl.ch/Publications/ById/350.html

[9] W. Chen, S. Toueg, and M. K. Aguilera, “On the quality of service of failure detectors,” *IEEE Trans. on Computers*, vol. 51, no. 2, pp. 561–580, May 2002.

[10] P. Urbán, I. Shnayderman, and A. Schiper, “Comparison of failure detectors and group membership: Performance study of two atomic broadcast algorithms,” in *Proc. Int’l Conf. on Dependable Systems and Networks*, San Francisco, CA, USA, June 2003, pp. 645–654. [Online]. Available: http://lsrwww.epfl.ch/Publications/ById/349.html

[11] N. Sergent, X. Défago, and A. Schiper, “Impact of a failure detection mechanism on the performance of consensus,” in *Proc. IEEE Pacific Rim Symp. on Dependable Computing (PRDC)*, Seoul, Korea, Dec. 2001, pp. 137–145. [Online]. Available: http://lsrwww.epfl.ch/Publications/ById/292.html

[12] A. Coccoli, P. Urbán, A. Bondavalli, and A. Schiper, “Performance analysis of a consensus algorithm combining Stochastic Activity Networks and measurements,” in *Proc. Int’l Performance and Dependability Symp.*, Washington, DC, USA, June 2002, pp. 551–560. [Online]. Available: http://lsrwww.epfl.ch/Publications/ById/294.html

[13] L. Sampaio, F. V. Brasileiro, W. d. C. Cirne, and J. de Figueiredo, “How bad are wrong suspicious: Towards adaptive distributed protocols,” in *Proc. Int’l Conf. on Dependable Systems and Networks (DSN)*, San Francisco, CA, USA, June 2003.

[14] N. Hayashibara, P. Urbán, A. Schiper, and T. Katayama, “Performance comparison between the Paxos and Chandra-Toueg consensus algorithms,” in *Proc. Int’l Arab Conf. on Information Technology (ACIT 2002)*, Doha, Qatar, Dec. 2002, pp. 526–533. [Online]. Available: http://lsrwww.epfl.ch/Publications/ById/346.html

[15] P. Urbán, X. Défago, and A. Schiper, “Chasing the FLP impossibility result in a LAN or how robust can a fault tolerant server be?” in *Proc. 20th IEEE Symp. on Reliable Distributed Systems (SRDS)*, New Orleans, LA, USA, Oct. 2001, pp. 190–193. [Online]. Available: http://lsrwww.epfl.ch/Publications/ById/288.html

[16] H. Ramasamy, P. Pandey, J. Lyons, M. Cukier, and W. Sanders, “Quantifying the cost of providing intrusion tolerance in group communication systems,” in *Proc. 2002 Int’l Conf. on Dependable Systems and Networks (DSN-2002)*, Washington, DC, USA, June 2002, pp. 229–238.

[17] T. D. Chandra, V. Hadzilacos, and S. Toueg, “The weakest failure detector for solving consensus,” *Journal of the ACM*, vol. 43, no. 4, pp. 685–722, July 1996. [Online]. Available: http://www.acm.org/pubs/toc/Abstracts/jacm/234549.html

[18] C. Fetzer, M. Raynal, and F. Tronel, “An adaptive failure detection protocol,” in *Proc. 8th IEEE Pacific Rim Symp. on Dependable Computing (PRDC-8)*, Seoul, Korea, Dec. 2001.

[19] V. Hadzilacos and S. Toueg, “A modular approach to fault-tolerant broadcasts and related problems,” Dept. of Computer Science, Cornell University, Ithaca, NY, USA, TR 94-1425, May 1994.

[20] A. Mostéfaoui and M. Raynal, “Solving consensus using Chandra-Toueg’s unreliable failure detectors: A general quorum-based approach,” in *Proc. 13th Int’l Symp. on Distributed Computing (DISC)*, no. 1693, Bratislava, Slovak Republic, Sept. 1999, pp. 49–63.

[21] S. Frolund and F. Pedone, “Revisiting reliable broadcast,” HP Laboratories, Palo Alto, CA, USA, Tech. Rep. HPL-2001-192, Aug. 2001. [Online]. Available: http://www.hpl.hp.com/techreports/2001/HPL-2001-192.html

[22] X. Défago and A. Schiper, “Specification of replication techniques, semi-passive replication, and lazy consensus,” Japan Advanced Institute of Science and Technology, Ishikawa, Japan, Research Report KS-RR-2002-001, Feb. 2002.

[23] E. Anceaume, “A lightweight solution to uniform atomic broadcast for asynchronous systems,” in *Proc. 27th Int’l Symp. on Fault-Tolerant Computing (FTCS-27)*, Seattle, WA, USA, June 1997, pp. 292–301.

[24] R. Boichat, P. Dutta, S. Frlund, and R. Guerraoui, “Deconstructing Paxos,” *SIGACT News*, vol. 34, no. 1, pp. 47–67, 2003.

[25] R. Friedman and R. van Renesse, “Packing messages as a tool for boosting the performance of total ordering protocols,” Dept. of Computer Science, Cornell University, Ithaca, NY, USA, TR 95-1527, July 1995.

[26] F. B. Schneider, “Implementing fault-tolerant services using the state machine approach: A tutorial,” *ACM Computing Surveys*, vol. 22, no. 4, pp. 299–319, Dec. 1990.

[27] H. Madeira, K. Kanoun, J. Arlat, Y. Crouzet, A. Johansson, R. Lindström, *et al.*, “Preliminarily dependability benchmark framework,” Dependability Benchmarking project (DBench), EC IST-2000-25425, Project deliverable CF2, Aug. 2001. [Online]. Available: http://www.laas.fr/DBench/

[28] P. Urbán, X. Défago, and A. Schiper, “Neko: A single environment to simulate and prototype distributed algorithms,” *Journal of Information Science and Engineering*, vol. 18, no. 6, pp. 981–997, Nov. 2002. [Online]. Available: http://lsrwww.epfl.ch/Publications/ById/307.html

[29] K. Tindell, A. Burns, and A. J. Wellings, “Analysis of hard real-time communications,” *Real-Time Systems*, vol. 9, no. 2, pp. 147–171, Sept. 1995.

[30] J. Gray, “Why do computers stop and what can be done about it ?” in *Proc. 5th Symp. on Reliablity in Distributed Software and Database systems*, Jan. 1986.

Appendix I
Full set of results for the normal-steady faultload

A. Graphs showing the early latency

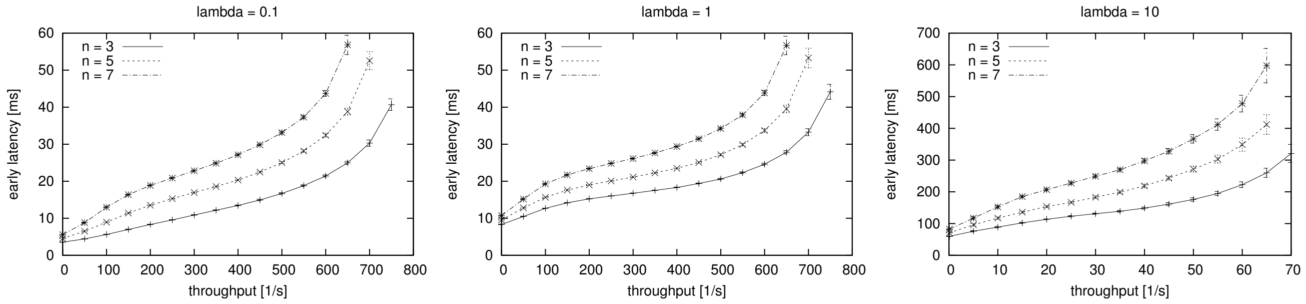

Fig. 11. Latency vs. throughput with the normal-steady faultload.

B. Graphs showing the late latency

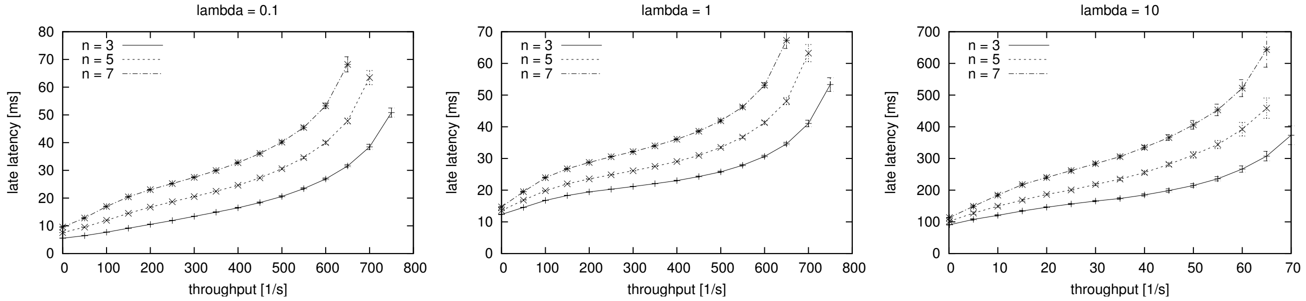

Fig. 12. Latency vs. throughput with the normal-steady faultload.

Appendix II
Full set of results for the crash-steady faultload

A. Graphs showing the early latency

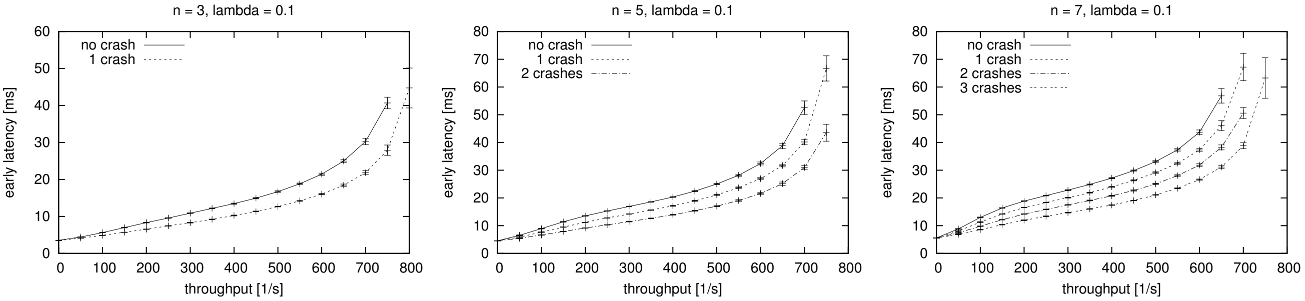

Fig. 13. Latency vs. throughput with the crash-steady faultload ($\lambda = 0.1$).

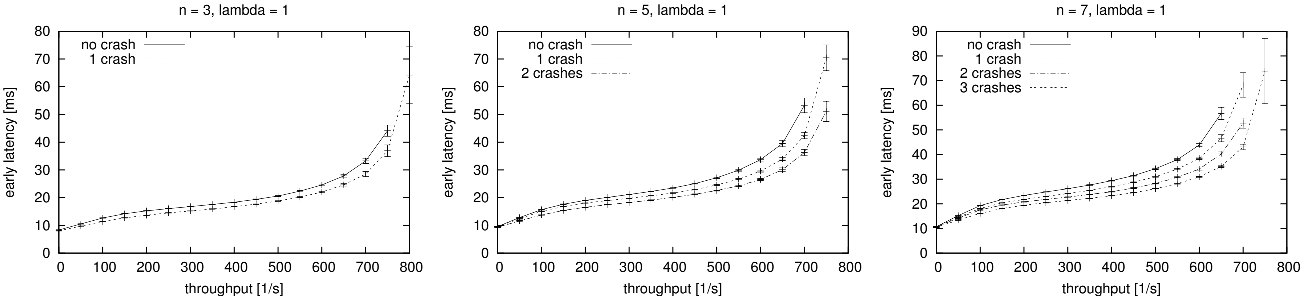

Fig. 14. Latency vs. throughput with the crash-steady faultload ($\lambda = 1$).

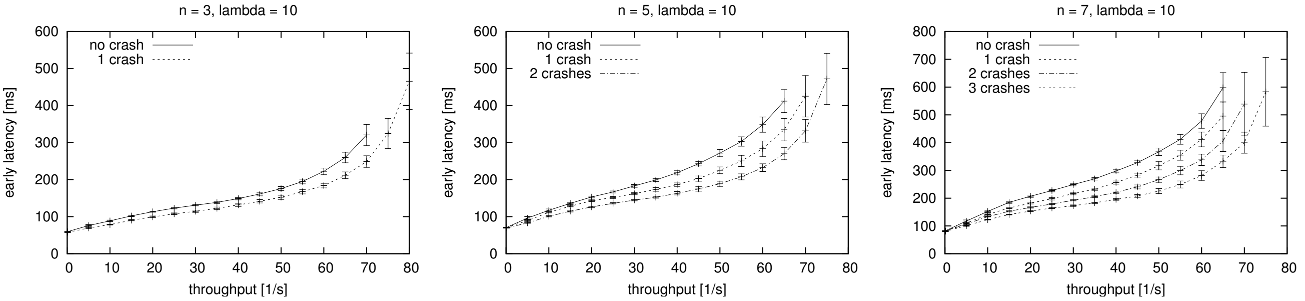

Fig. 15. Latency vs. throughput with the crash-steady faultload ($\lambda = 10$).

B. *Graphs showing the late latency*

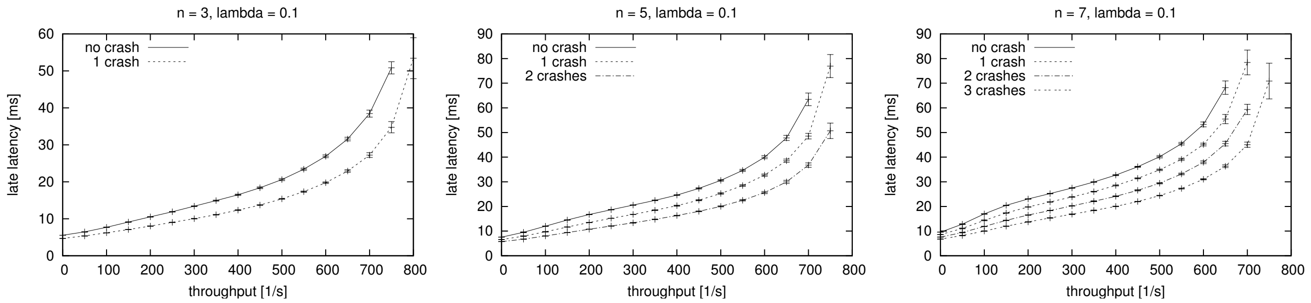

Fig. 16. Latency vs. throughput with the crash-steady faultload ($\lambda = 0.1$).

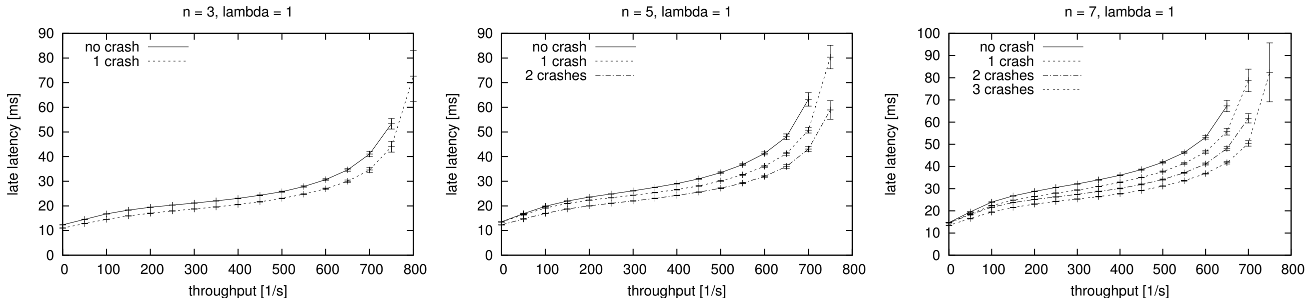

Fig. 17. Latency vs. throughput with the crash-steady faultload ($\lambda = 1$).

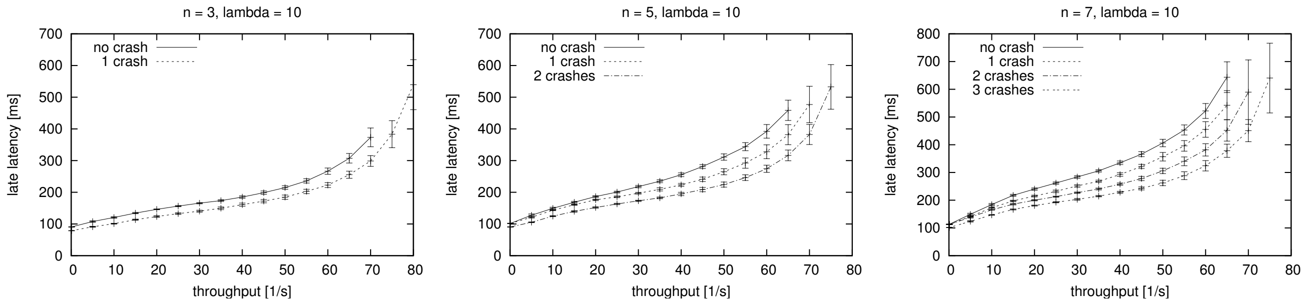

Fig. 18. Latency vs. throughput with the crash-steady faultload ($\lambda = 10$).

APPENDIX III
FULL SET OF RESULTS FOR THE CRASH-TRANSIENT FAULTLOAD

A. Graphs showing the early latency overhead

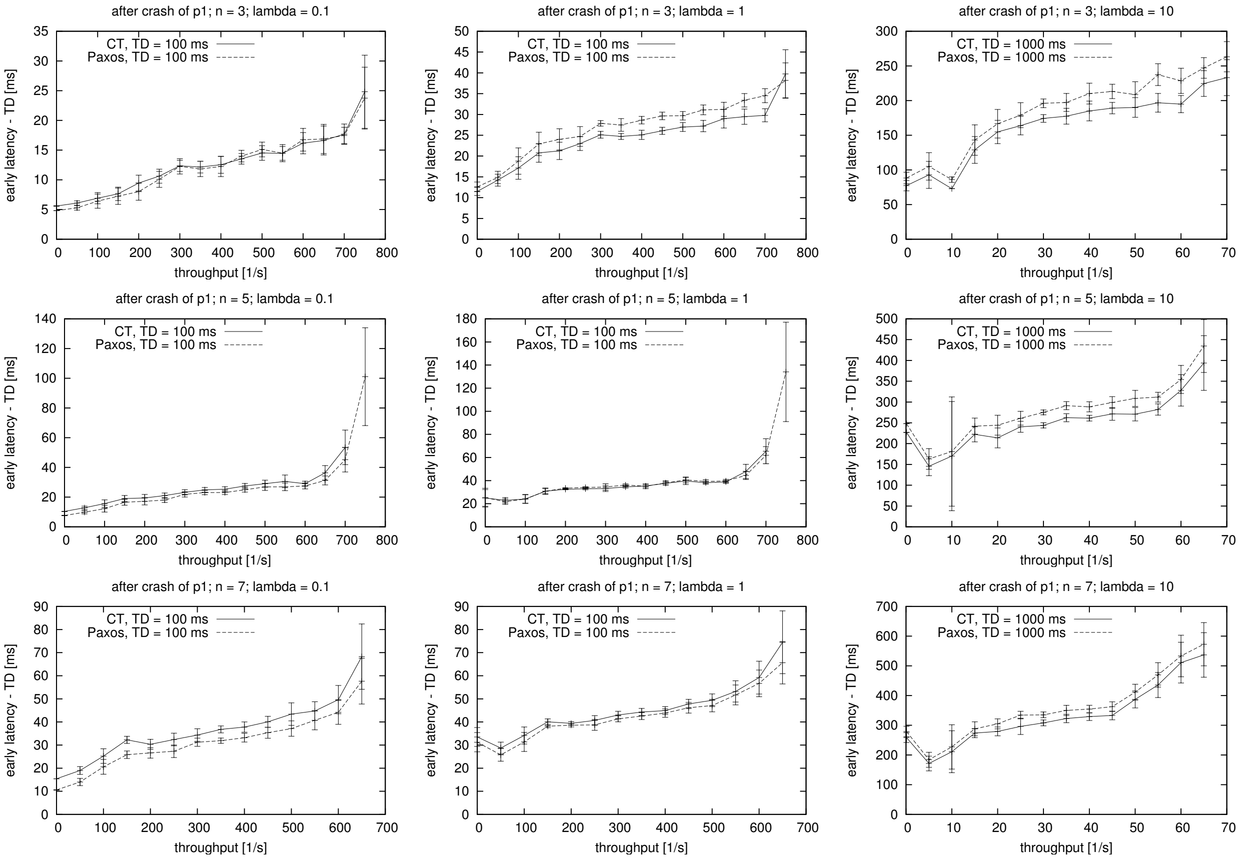

Fig. 19. Latency overhead vs. throughput with the crash-transient faultload. One process crashes.

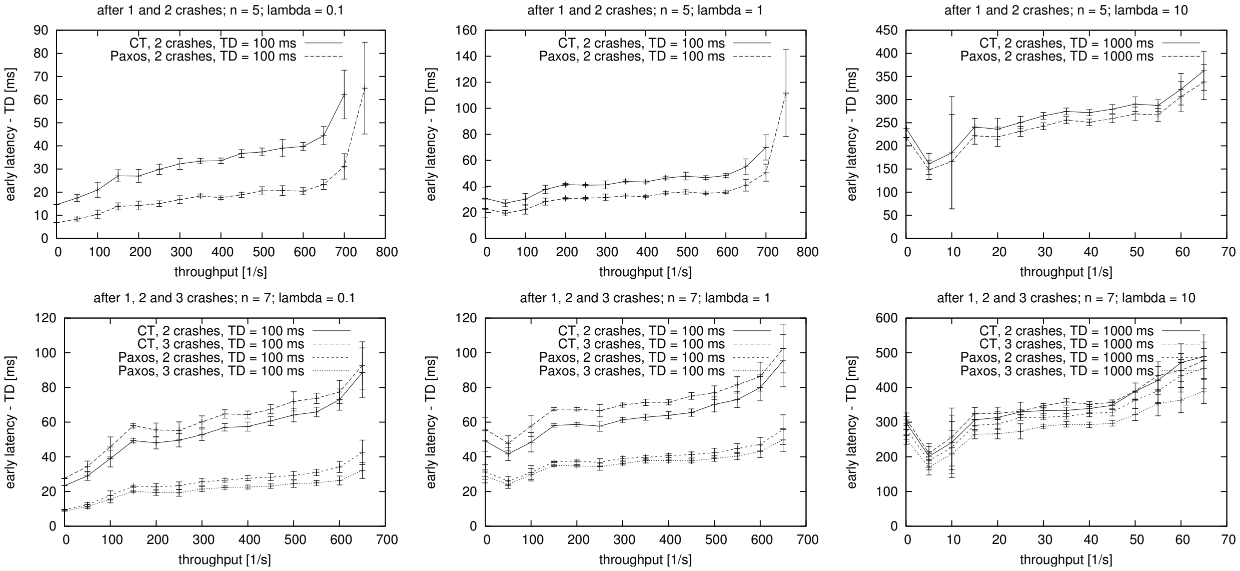

Fig. 20. Latency overhead vs. throughput with the crash-transient faultload. Multiple processes crash.

B. Graphs showing the late latency overhead

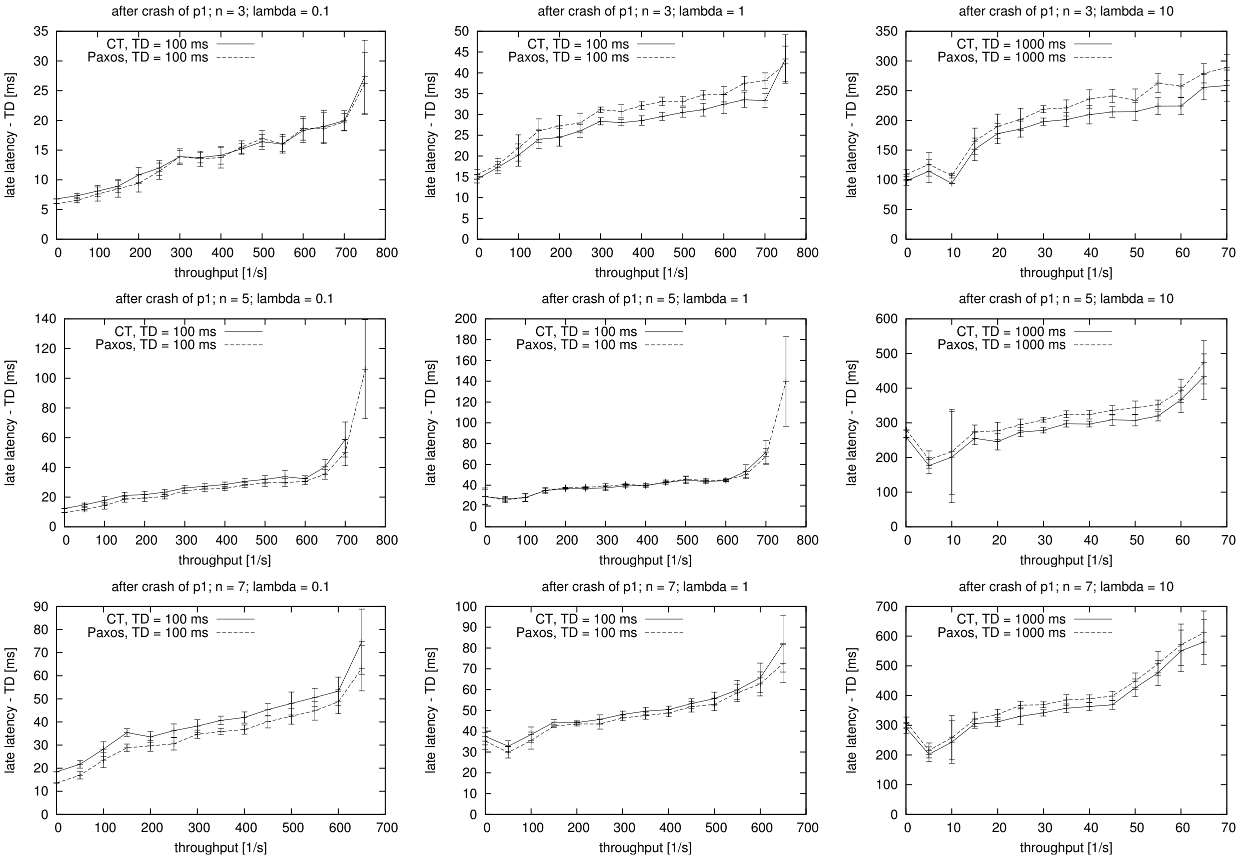

Fig. 21. Latency overhead vs. throughput with the crash-transient faultload. One process crashes.

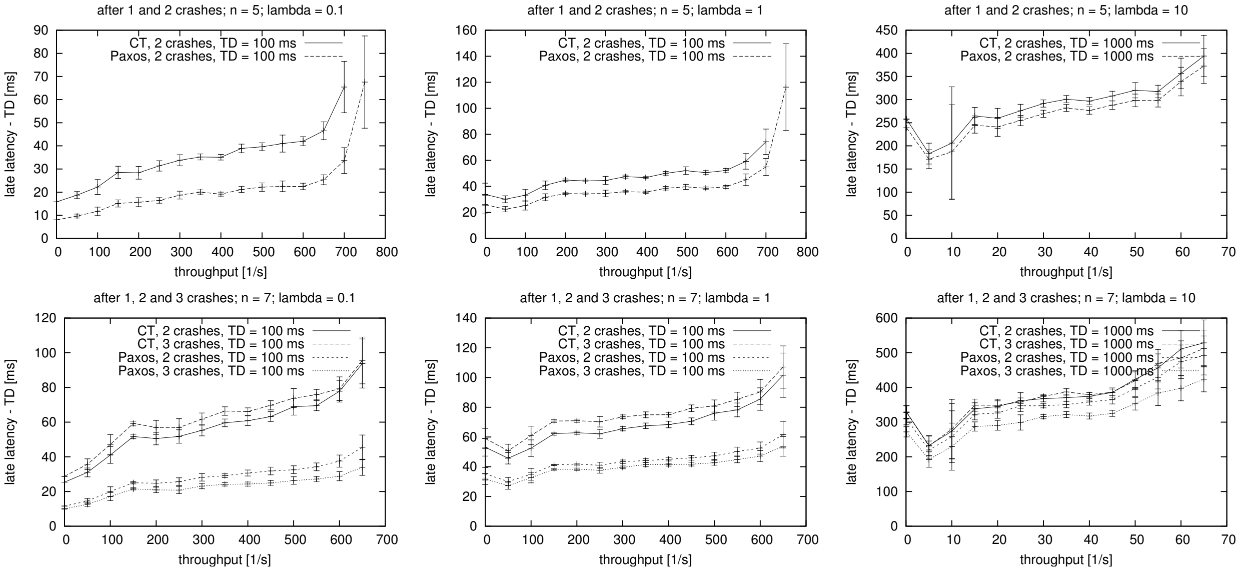

Fig. 22. Latency overhead vs. throughput with the crash-transient faultload. Multiple processes crash.

APPENDIX IV
FULL SET OF RESULTS FOR THE SUSPICION-STEADY FAULTLOAD

A. Graphs showing the early latency

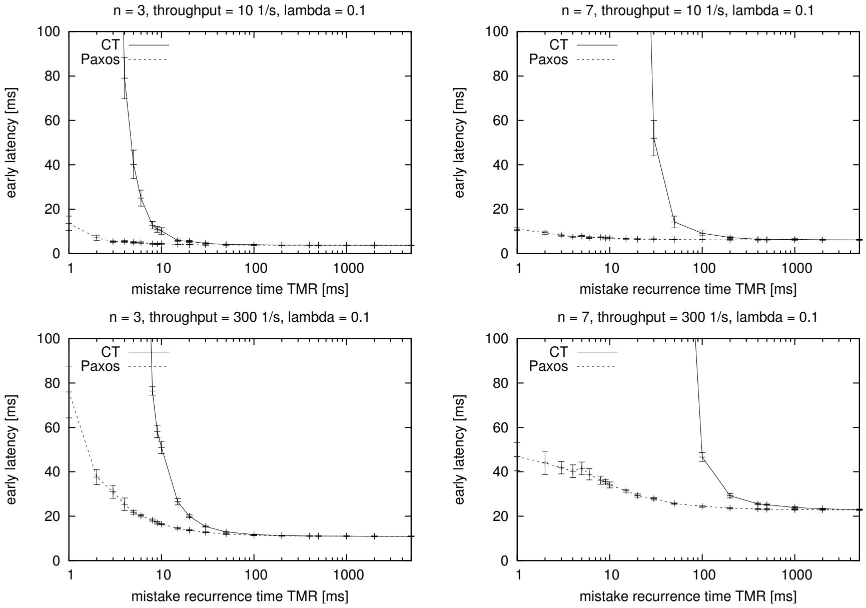

Fig. 23. Latency vs. T_{MR} with the suspicion-steady faultload, with $T_M = 0$ ($\lambda = 0.1$).

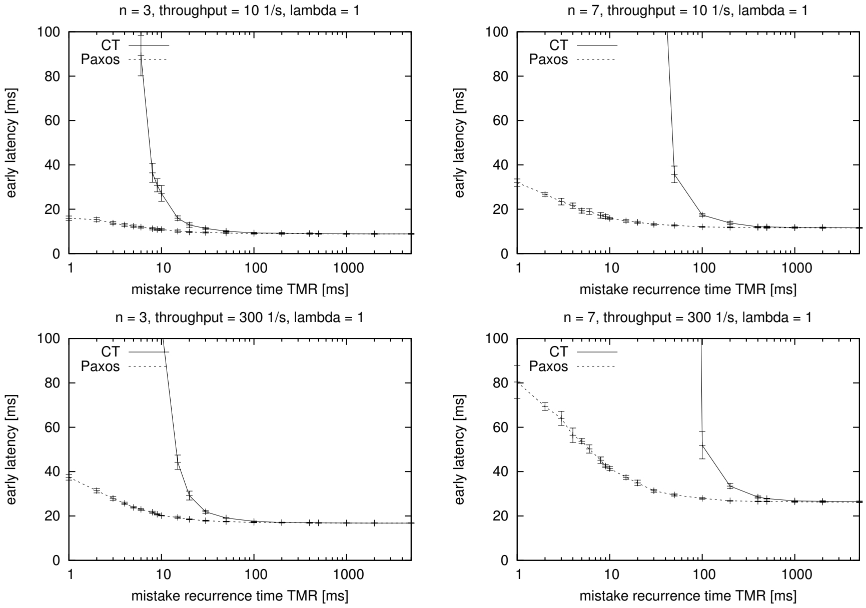

Fig. 24. Latency vs. T_{MR} with the suspicion-steady faultload, with $T_M = 0$ ($\lambda = 1$).

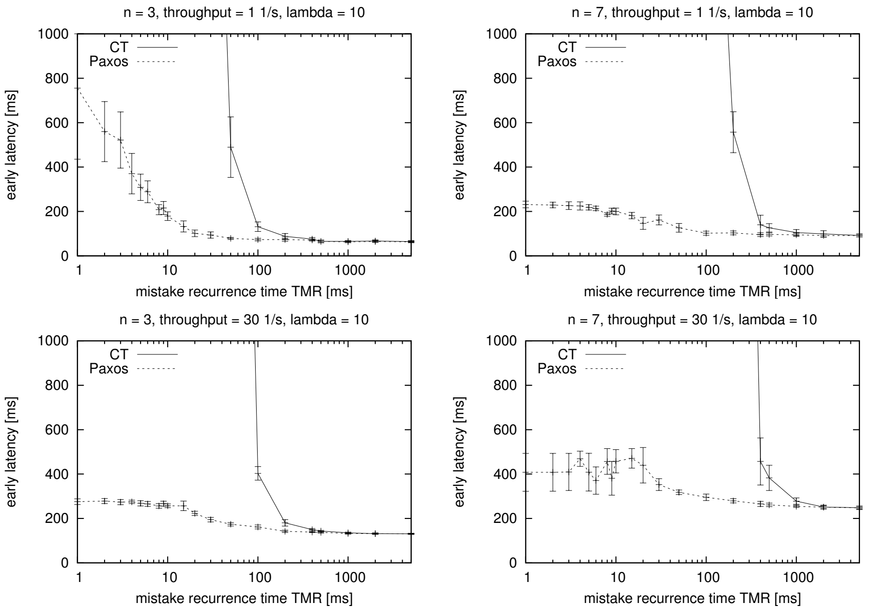

Fig. 25. Latency vs. T_{MR} with the suspicion-steady faultload, with $T_M = 0$ ($\lambda = 10$).

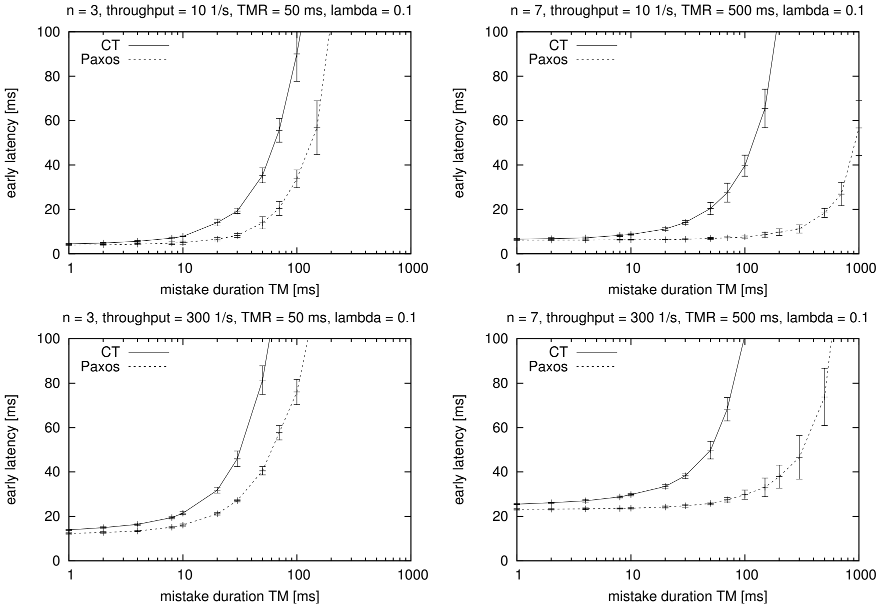

Fig. 26. Latency vs. T_M with the suspicion-steady faultload, with T_{MR} fixed ($\lambda = 0.1$).

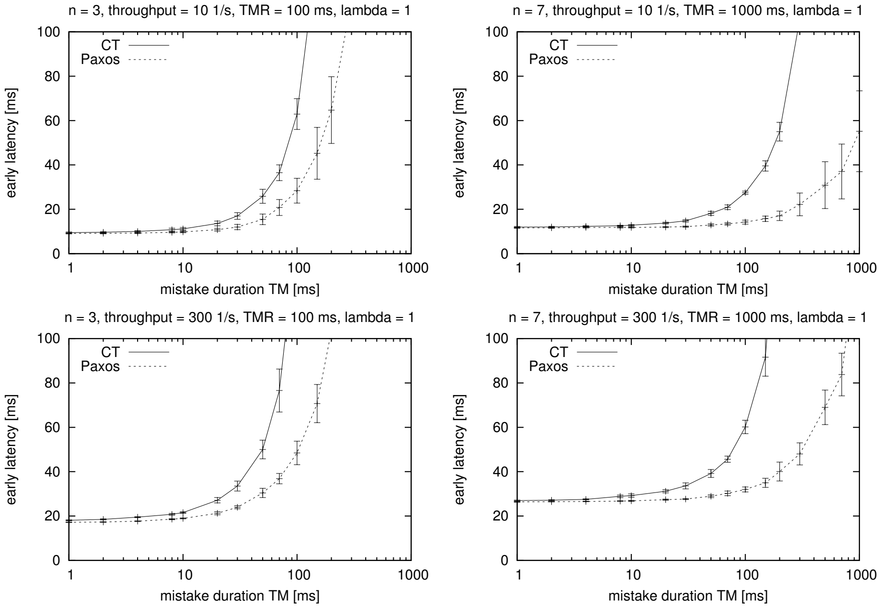

Fig. 27. Latency vs. T_M with the suspicion-steady faultload, with T_{MR} fixed ($\lambda = 1$).

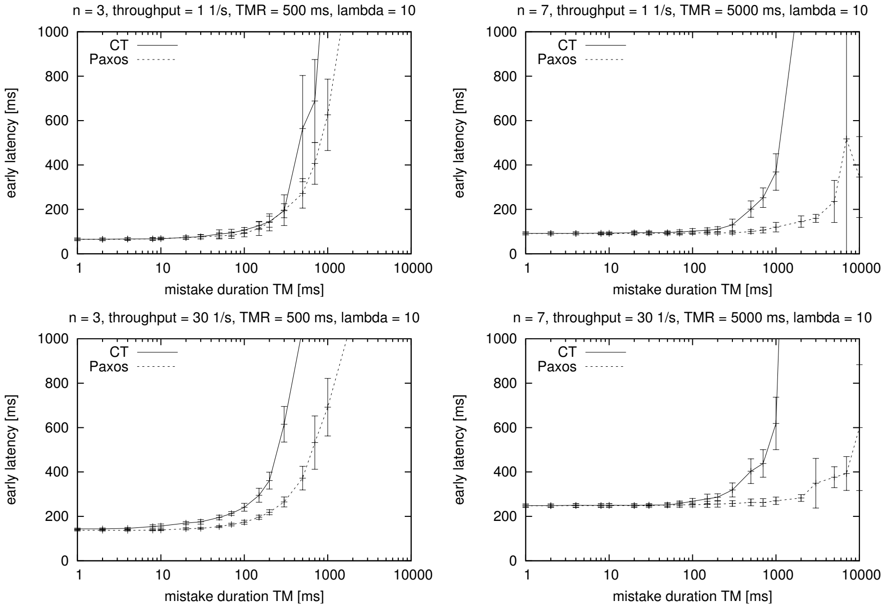

Fig. 28. Latency vs. T_M with the suspicion-steady faultload, with T_{MR} fixed ($\lambda = 10$).

B. Graphs showing the late latency

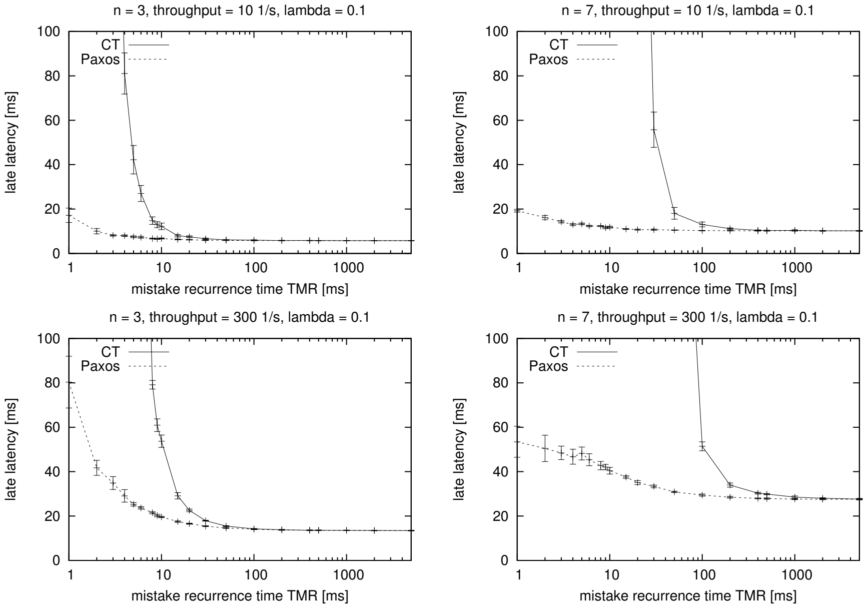

Fig. 29. Latency vs. T_{MR} with the suspicion-steady faultload, with $T_M = 0$ ($\lambda = 0.1$).

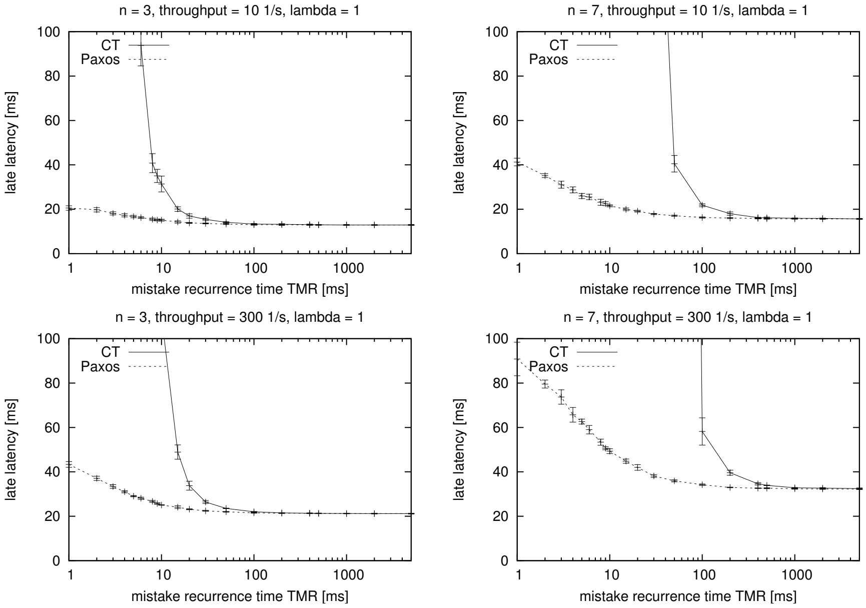

Fig. 30. Latency vs. T_{MR} with the suspicion-steady faultload, with $T_M = 0$ ($\lambda = 1$).

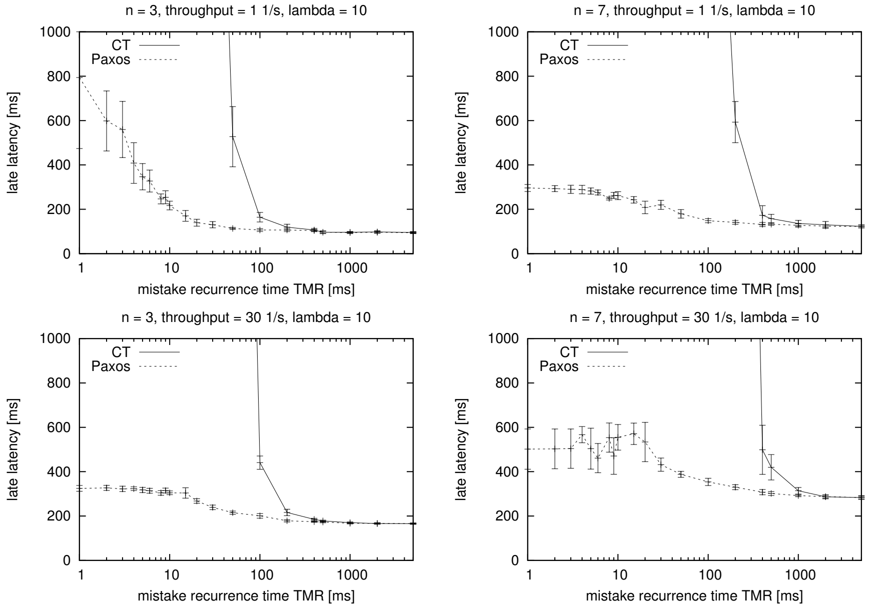

Fig. 31. Latency vs. T_{MR} with the suspicion-steady faultload, with $T_M = 0$ ($\lambda = 10$).

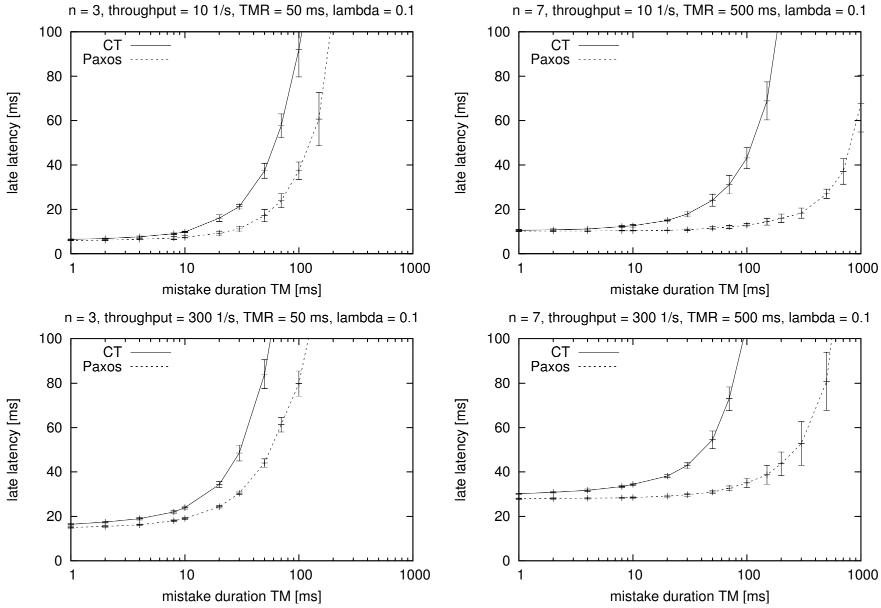

Fig. 32. Latency vs. T_M with the suspicion-steady faultload, with T_{MR} fixed ($\lambda = 0.1$).

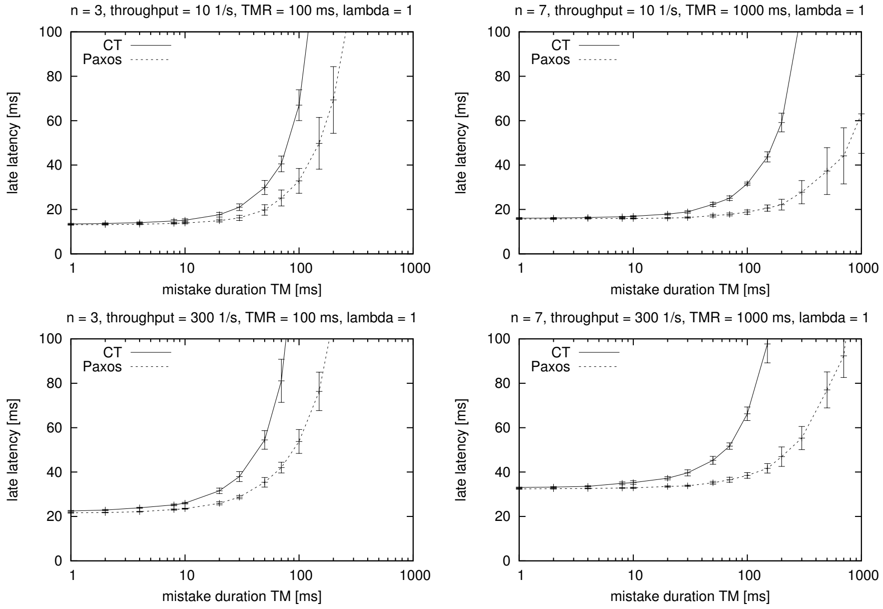

Fig. 33. Latency vs. T_M with the suspicion-steady faultload, with T_{MR} fixed ($\lambda = 1$).

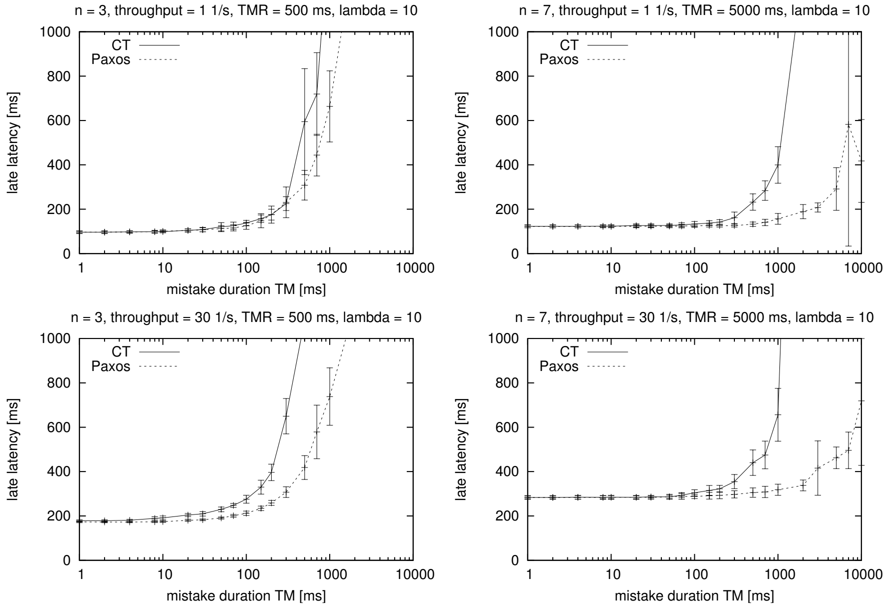

Fig. 34. Latency vs. T_M with the suspicion-steady faultload, with T_{MR} fixed ($\lambda = 10$).