Proc. 22nd Symposium on Reliable Distributed Systems (SRDS'03), Florence, Italy, 361–370 November 13, 2003

Randomized Asynchronous Consensus with Imperfect Communications

Ulrich Schmid*
Technische Universität Wien
Embedded Computing Systems Group (E182/2)
Treitlstraße 3, A-1040 Vienna
s@ecs.tuwien.ac.at

Christof Fetzer
AT&T Labs-Research
Dependable Distributed Computing Department
180 Park Ave, Florham Park, NJ07932
christof@research.att.com

Abstract

We introduce a novel hybrid failure model, which facilitates an accurate and detailed analysis of round-based synchronous, partially synchronous and asynchronous distributed algorithms under both process and link failures. Granting every process in the system up to f_ℓ send and receive link failures (with f_ℓ^a arbitrary faulty ones among those) in every round, without being considered faulty, we show that the well-known randomized Byzantine agreement algorithm of (Srikanth & Toueg 1987) needs just $n \geq 4f_\ell + 2f_\ell^a + 3f_a + 1$ processes for coping with f_a Byzantine faulty processes. The probability of disagreement after R iterations is only 2^{-R}, which is the same as in the FLP model and thus much smaller than the lower bound $\mathcal{O}(1/R)$ known for synchronous systems with lossy links. Moreover, we show that 2-stubborn links are sufficient for this algorithm. Hence, contrasting widespread belief, a perfect communications subsystem is not required for efficiently solving randomized Byzantine agreement.

1 Motivation

Most asynchronous fault-tolerant algorithms have been designed for process failures only. A common assumption is that at most f of the n processes in a distributed system may be Byzantine faulty during the entire execution. Link failures can also be handled within such models, as long as link failure rates are low. If link failure rates are high, however, one cannot count link-related failures as sender or receiver process failures without quickly running out of non-faulty processes [22, 23].

Most failure models for asynchronous systems hence stipulate [11] or simulate [1, 2, 4, 5] perfect links, i.e., can assume that any message sent by a correct process to any other correct process will eventually be delivered. Although many distributed protocols have been designed for reliable channels, this abstraction inevitably needs unbounded memory and stable storage or, equivalently, UIDs in presence of crashes [10, 18]. Consequently, in wide area and wireless networks, where long lasting link failures and non-FIFO behavior are common, the perfect communications assumption of the Fischer, Lynch & Paterson (FLP) model [11] is difficult to justify.

As the first contribution of this paper, we introduce a hybrid failure model that incorporates both process and link failures in both time and value domain. It is applicable to any round-based distributed algorithm and consists of a basic *physical failure model* facilitating assumption coverage and timing analysis, and a more abstract round-by-round *perception failure model* facilitating accurate fault-tolerance analysis. Formalized for partially synchronous systems [9] where delays are possibly unknown, our model covers synchronous and asynchronous systems as limiting cases and is hence widely applicable.

We will demonstrate the utility of our failure model by analyzing a hybrid version of the randomized Byzantine agreement algorithm of [25, 26] under Byzantine process and link failures. For synchronous systems, it has long been known that randomized algorithms—unlike deterministic ones, see [22, 23]—can achieve consensus in presence of (unconstrained) lossy links [18]. There is, however, a fairly large lower bound $1/(R+1)$ for the probability of disagreement after R rounds [18, Thm. 5.5]. For asynchronous systems under the FLP model, the algorithm of [26] has a probability of disagreement of at most 2^{-R}. The question is: Does this algorithm still work if one drops the perfect link assumption and, if so, is there a penalty?

In the following sections[1], we will show the following: In Section 3, we demonstrate how to adapt two well-known reliable broadcast primitives to the failure model of Sec-

*Supported by the Austrian START programme Y41-MAT, in the context of our W2F-project *http://www.auto.tuwien.ac.at/Projects/W2F/*

[1] Due to lack of space, we can only provide some of our results and derivations here; consult [21] for the full version of our paper.

tion 2. They are the major building blocks of the randomized consensus algorithm of [25, 26], which is analyzed in Section 4. Our results reveal that a large number of link failures can be tolerated system-wide, without increasing the probability of disagreement 2^{-R}, provided that the number of processes in the system is moderately increased. In Section 5, we finally show that 2-stubborn links [14] (guaranteeing reliable delivery of the last two messages) can be used for even higher link failure rates. A short summary of our accomplishments in Section 6 concludes our paper.

2 Failure Model

This section contains a very brief overview of our failure model. It consists of an execution model, a basic physical failure model, and a more abstract perception failure model. Both the physical and the perception failure model are *hybrid* ones [3,27], i.e., distinguish several classes of failures. The advantage of a hybrid failure model is its improved resilience: Less severe failures can usually be handled with fewer processes than more severe ones. Obviously, an algorithm's resilience in a standard model (like all-Byzantine) is easily obtained by setting some model parameters to 0.

Due to lack of space, we will entirely omit the description of the *physical failure model*, which is an extension of the model of [20]. It distinguishes several classes of time and value failures for both processes and links, and uses assertions like "at most ϕ_{av} processes may be Byzantine". Due to the exploding number of possible combinations of time and value failures, it is, however, not used for analyzing fault-tolerant algorithms. Its primary purpose is the analysis of the assumption coverage [16] in real systems.

The physical failure model can be reduced to a more abstract (and vastly simpler) *perception failure model*, which is similar in spirit to the round-by-round fault detector approach of [12]. It is a generalization of the synchronous model of [22,23], and is solely based upon the local view (= perception of failures) of every process in a round. The perception failure model is particularly well-suited for analyzing the fault-tolerance properties of distributed algorithms.

The Execution Model

We consider a distributed system of n *processors* connected by a fully or partially connected point-to-point network. All links between processors are bidirectional, consisting of two unidirectional channels that may be hit by failures independently. The system will execute a distributed round-based algorithm made up of one or more concurrent *processes* at every processor. Any two processes at different processors can communicate bidirectionally with each other via the interconnecting links. Every processor is identified by a unique *processor id* $p \in \{1, \ldots, n\}$; every process is uniquely identified system-wide by the tuple (processor id, process name), where the *process name* N is chosen from a suitable name space. We will distinguish processes primarily by their processor ids, however, i.e., suppress process names when they are clear from the context.

Since we restrict our attention to round-based algorithms, all processes execute a finite or infinite sequence of consecutive *rounds* $R = 0, 1, \ldots$. In every round—except the initial one $R = 0$, which is slightly different—a single process p may broadcast (= successively send) a single message containing the current *round number* R and a *value* V_p^R depending upon its local computation to all processes contained in p's current *receiver set* $\mathcal{R}_p^R \subseteq \{1, \ldots, n\}$. We assume that every (non-faulty) receiver q knows its current *sender set* $\mathcal{S}_q^R = \{p : q \in \mathcal{R}_p^R\}$ containing all the processes that should have sent a message to it, and that a process satisfying $p \in \mathcal{R}_p^R$ (and hence $p \in \mathcal{S}_p^R$) sends a message to itself as well. Note that this convention does not prohibit an efficient direct implementation of self-reception, provided that the resulting end-to-end transmission delay is taken into account properly.

Concurrently, for every round number S, process p receives incoming round S messages from the processes $\in \mathcal{S}_p^S$ and collects their values in a local array (subsequently called *perception vector*) $\mathcal{V}_p^S = \{V_p^{1,S}, \ldots, V_p^{n,S}\}$. Note that the perception vector as well as its individual entries are time-dependent, i.e., $\mathcal{V}_p^S = \mathcal{V}_p^S(t)$ and $V_p^{i,S} = V_p^{n,S}(t)$. However, we will usually suppress t in order not to overload our notation. Storing a single value for each peer process in the perception vector is sufficient, since any receiver may get at most one round S message from any non-faulty sender. The entry $V_p^{q,S} \in \mathcal{V}_p^S$ (subsequently called *perception*) is either $\emptyset$ if no round S message from process q came in yet (or if $q \notin \mathcal{S}_p^S$), or it contains the received value from the first round S message from process q. In case of multiple messages from the same sender q, which implies that q is faulty, the receiver could also drop all messages and set $V_p^{q,S}$ to some obviously faulty value, instead of retaining the value from the first message.

Process p's current round R is eventually terminated at the *round switching time* σ_p^R, which is the real-time when process p switches from round R to the next round $R+1$. The time interval of p's round $R+1$ is hence $(\sigma_p^R, \sigma_p^{R+1}]$. In asynchronous systems, round switching is event-based and part of the particular algorithm. In synchronous systems, round switching is enforced externally. At the round switching time, the value $V_p^{R+1} = F_p\big(\mathcal{V}_p^R(\sigma_p^R), \Sigma_p\big)$ to be broadcast by process p in the next round $R+1$ is computed as a function F_p of the round R perceptions available in $\mathcal{V}_p^R = \mathcal{V}_p^R(\sigma_p^R)$ and p's local state Σ_p at time σ_p^R.

Formally, the essentials of the above execution pattern are captured by two specific events: $be_p^R = V_p^R(t_p^R)$ is process p's round R *broadcast event*, whereas $pe_q^{p,R} =$

$V_q^{p,R}(t_q^{p,R})$ denotes process q's *perception event* of process p's broadcast event. Those events are related via their parameter values $V_q^{p,R} = V_p^R$ (which are equal if there is no failure) and their occurrence times $t_q^{p,R} = t_p^R + \delta_q^{p,R}$, where $\delta_q^{p,R}$ is the *end-to-end computational + transmission delay* between sender p and receiver q in round R. Note that $\delta_q^{p,R}$ includes any round R computation at the sender and receiver process, in particular, $F_p\big(\mathcal{V}_p^R(\sigma_p^R), \Sigma_p\big)$.

Our model stipulates lower and upper bounds $\tau^- \geq 0$ and $\tau^+ \leq \infty$, not necessarily known to the algorithm, such that[2] $\tau^- \leq \delta_q^{p,R} \leq \tau^+$ for any two well-behaved processes p, q connected by a non-faulty link. Note that this relation must be valid for any round R and for $p = q$ as well.[3] Introducing the interval $\boldsymbol{\tau} = [\tau^-, \tau^+]$, this requirement can be written concisely as $\delta_q^{p,R} \in \boldsymbol{\tau}$. The resulting bound for $\delta_q^{p,R}$'s *uncertainty*, which will play a central role in our perception failure model, is given by $\varepsilon = \tau^+ - \tau^-$.

The Perception Failure Model

Consider the round R perception vector $\mathcal{V}_q^R(t)$— observed at some real-time t— of a well-behaved process q. First of all, our execution model implies that $\mathcal{V}_q^R(t)$ is monotonic in time, in the sense that $|\mathcal{V}_q^R(t+\Delta t)| \geq |\mathcal{V}_q^R(t)|$ for any $\Delta t \geq 0$, since perceptions are only added. Moreover, since the value V_q^{R+1} to be broadcast in the next round $R+1$ is computed solely from $\mathcal{V}_q^R := \mathcal{V}_q^R(\sigma_q^R)$ and q's local state at the round switching time σ_q^R, it is obvious that, ultimately, only the failures in the perceptions present at the respective round switching times count. Timing failures are no longer visible here (but will probably affect σ_q^R), since a message that was not received by σ_q^R at process q just results in $V_q^{p,R} = \emptyset$. Consequently, the resulting *perception failure model* is much simpler than the physical one and therefore more suitable for analyzing an algorithm's fault-tolerance properties.

Our formalization solely rests upon the $n \times n$ matrix $\mathcal{V}(t) = \mathcal{V}^R(t)$ of round R perceptions observed at the same arbitrary time t—typically, some process's round switching time—at all processes:

$$\mathcal{V}(t) = \begin{pmatrix} \mathcal{V}_1(t) \\ \mathcal{V}_2(t) \\ \vdots \\ \mathcal{V}_n(t) \end{pmatrix} = \begin{pmatrix} V_1^1 & V_1^2 & \dots & V_1^n \\ V_2^1 & V_2^2 & \dots & V_2^n \\ \vdots & \vdots & \vdots & \vdots \\ V_n^1 & V_n^2 & \dots & V_n^n \end{pmatrix}_t \quad (1)$$

Note that $\mathcal{V}(t)$ is in fact a quite flexible basis for our failure model, since different "views" of the state of the distributed computation can be produced easily by choosing a suitable time of observation t.

We distinguish the following failure modes for single perceptions in $\mathcal{V}(t)$ in our perception failure model:

Definition 1 (Perception Failures) *Process q's perception V_q^p of process p's broadcast value V_p can be classified according to the following mutually exclusive failure mode predicates:* $\text{correct}\big(V_q^p\big)$*:* $V_q^p = V_p$*;* $\text{omission}\big(V_q^p\big)$*:* $V_q^p = \emptyset$*;* $\text{value}\big(V_q^p\big)$*:* $V_q^p \neq \emptyset$ *and* $V_q^p \neq V_p$*.*

Next, we have to classify sender[4] process failures. This requires the important notion of *obedient* processes: An obedient process is an alive process that faithfully executes the particular algorithm. It gets its inputs and performs its computations exactly as a non-faulty process, but it might fail in specific ways to communicate its value to the outside world. We will subsequently use this term instead of non-faulty whenever a process acts as a receiver ("obedient receiver"), since this will allow us to reason about the behavior of (benign) faulty processes as well. If $\mathcal{R}_p$ denotes some sender p's receiver set, let $\overline{\mathcal{R}}_p \subseteq \mathcal{R}_p$ denote the set of obedient processes among those.

Whereas the physical failure model differentiates timing failures according to $\delta_q^{p,R} \in \boldsymbol{\tau}$ vs. $\delta_q^{p,R} \notin \boldsymbol{\tau}$ and hence incorporates those quantities explicitly, it is solely the choice of t that is used in Definition 2 for this purpose: It only depends upon t whether a non-faulty perception V_q^p represents a perception event pe_q^p from a non-faulty or rather a timing faulty process p. Hence, neither $\delta_q^{p,R}$ nor $\boldsymbol{\tau}$ will show up in the definitions of the perception failure model below.

Definition 2 (Perception Process Failures) *Let p be a (faulty) sender process and g be some obedient receiver with $\emptyset \neq V_g^p \in \mathcal{V}(t)$, if there is any such g. In the absence of link failures, process failures of p can be classified according to the perceptions $V_q^p \in \mathcal{V}(t+\varepsilon)$ at all obedient receivers $q \in \overline{\mathcal{R}}_p \subseteq \mathcal{R}_p$ as follows:*

- *Non-faulty:* $\forall q \in \overline{\mathcal{R}}_p : \text{correct}\big(V_q^p\big)$,
- *Manifest:* $\forall q, r \in \overline{\mathcal{R}}_p : V_q^p = V_r^p \neq V_p$ *detectably,*
- *Clean crash:* $\forall q \in \overline{\mathcal{R}}_p : \text{omission}\big(V_q^p\big)$,
- *Omission:* $\forall q \in \overline{\mathcal{R}}_p : \text{correct}\big(V_q^p\big) \vee \text{omission}\big(V_q^p\big)$,
- *Symmetric:* $\forall q, r \in \overline{\mathcal{R}}_p : V_q^p = V_r^p$,
- *Arbitrary: no constraints.*

A faulty process producing at most omission failures is called benign faulty and is assumed to be obedient.

[2] Note that $\delta_q^{p,R} \leq \tau^+$ must be replaced by $\delta_q^{p,R} < \tau^+$ in case of $\tau^+ = \infty$. This conveniently models purely asynchronous systems, where only assertions like "eventually, ..." are possible.

[3] Those assumptions could be somewhat relaxed. In case of time-varying delays, we could, for example, allow τ^- and τ^+ to be different in different rounds. This is exploited in the Θ-model of [17]. The very small self-reception delay could be considered as an (early) link timing failure and hence be masked by increasing f_ℓ^{ra} and f_ℓ^{sa} by 1.

[4] Receiver process failures (receive omissions) will be considered below, when introducing link failures.

Bear in mind that both arbitrary and symmetric faulty processes, but not benign faulty ones, may also be faulty in the time domain.

Definition 3 (Perception Link Failures) *In the absence of sender process failures, a failure of the link from sender p to an obedient receiver q can be classified according to its effect upon q's perception $V_q^p \in \mathcal{V}(t)$ as follows: Link non-faulty: $V_q^p = V_p$; link omission: $V_q^p = \emptyset$; link arbitrary: no constraint. The failure classes up to link omission failures are called benign.*

To overcome the impossibility of consensus in presence of unrestricted link failures [13, 18], it turned out that send and receive link failures should be considered separately [22, 23]:

(A1^s) *Broadcast link failures*: For any single sender s, there are at most f_ℓ^s receivers q with a perception vector $\mathcal{V}_q$ that contains a faulty perception V_q^s from s.

(A1^r) *Receive link failures*: In any single process q's perception vector $\mathcal{V}_q$, there are at most f_ℓ^r faulty V_q^p.

Separating broadcast and receive link failures as above makes sense due to the fact that we consider the unidirectional channels, rather than the bidirectional links, as single fault containment regions: Broadcast link failures affect outbound channels, whereas receive link failures affect inbound channels. Still, broadcast and receive link failures are of course not independent of each other: If a message from process p to q is hit by a failure in p's message broadcast, it obviously contributes a failure in process q's message reception as well. Nevertheless, our failure model will consider (A1^s) and (A1^r) as independent of each other and of process failures, for any process in the system and any round. Only the model parameters f_ℓ^s and f_ℓ^r cannot be independently chosen (without restricting the link failure patterns), since the system-wide number of send and receive link failures must of course match. Hence, $f_\ell^s = f_\ell^r = f_\ell$ is the most natural choice.

Note that we allow every process in the system to commit up to f_ℓ^s broadcast and up to f_ℓ^r receive link failures, in every round, without considering the process as faulty in the usual sense. In addition, the particular links actually hit by a link failure may be different in different rounds. A process must be considered (omission) faulty, however, if it exceeds its budget f_ℓ^s of broadcast link failures in some round. Note that a process that experiences more than f_ℓ^r receive link failures in some round must usually be considered (arbitrary) faulty, since it might be unable to correctly follow the algorithm after such an event.

The following Definition 4 contains our complete perception-based failure model, which just specifies the properties of any round's perception matrix $\mathcal{V}^R(t)$. Note that this definition is valid for arbitrary times t, including those where the perceptions from some senders did not yet arrive (and are hence $\emptyset$).

Definition 4 (Asynchronous Perception Failure Model) *Let $\mathcal{V}^R(t)$ be the round R perception matrix of an asynchronous system of processes running on different processors that comply to our execution model. For any obedient receiver q, it is guaranteed that $V_q^{p,R} = \emptyset$ if $p \notin \mathcal{S}_q^R$ or if $V_q^{p,R}$ was not received by time t. Moreover:*

- *(P1) There are at most f_a, f_s, f_o, f_c, and f_m columns in $\mathcal{V}^R(t)$ that correspond to arbitrary, symmetric, omission, clean crash, and manifest faulty processes and may hence contain perceptions $V_q^{p,R}$ according to Definition 2.*
- *(A1^s) In every single column p, at most f_ℓ^s perceptions $V_q^{p,R} \in \mathcal{V}^R(t)$ corresponding to obedient receivers $q \in \overline{\mathcal{R}}_p^R \subseteq \mathcal{R}_p^R$ may differ from the ones obtained in the absence of broadcast link failures. At most $f_\ell^{sa} \leq f_\ell^s$ of those perceptions may be link arbitrary faulty.*
- *(A1^r) In every single row q corresponding to an obedient receiver, at most f_ℓ^r of the perceptions $V_q^{p,R} \in \mathcal{V}^R(t)$ corresponding to senders $p \in \mathcal{S}_q^R$ may differ from the ones obtained in the absence of receive link failures. At most $f_\ell^{ra} \leq f_\ell^r$ of those may be link arbitrary faulty.*
- *(A2) Process q can be sure about the origin p of $V_q^{p,R} \in \mathcal{V}^R(t)$.*
- *(A3) $\emptyset \neq V_q^{p,R} \in \mathcal{V}^R(t) \Rightarrow \emptyset \neq V_r^{p,R} \in \mathcal{V}^R(t+\varepsilon)$ for every non-faulty sender p connected to obedient receivers q and r via non-faulty links.*

Note that the effects of process failures (P1) and link failures (A1^s), (A1^r) are considered orthogonal; it can hence happen that a link failure hits a perception originating from a faulty sender process. This is also true for manifest and clean crash failures, where link arbitrary failures could create non-empty perceptions at some receivers.

The primary way of using our failure model in the analysis of agreement-type algorithms is the following: Given the perception vector $\mathcal{V}_q^R(\sigma_q^R)$ of some specific obedient receiver process q at its round switching time σ_q^R, it allows to determine how many perceptions will at least be present in any other process r's perception vector $\mathcal{V}_r^R(\sigma_q^R + \varepsilon)$ shortly thereafter. The following Lemma 1 formalizes this fact.

Lemma 1 (Difference in Perceptions) *At any time t, the perception vector $\mathcal{V}_q(t)$ of any process at an obedient receiver q may contain at most $f_\ell^{ra} + f_a + f_s$ timing/value-faulty perceptions $V_q^p \neq \emptyset$. Moreover, at most $\Delta f = f_\ell^{ra} + f_\ell^r + f_a + f_o$ perceptions V_r^p corresponding to $V_q^p \neq \emptyset$ may be missing in any other obedient receiver's $\mathcal{V}_r(t + \Delta t)$ for any $\Delta t \geq \varepsilon$.*

Proof: The first statement is an obvious consequence of Definition 4. To prove the second one, we note that at most $f_\ell^{ra} + f_a + f_o$ perceptions may have been available (partly too early) at q without being avaliable yet at r, additional $f_\ell^{ra\prime} \leq f_\ell^{ra}$ perceptions may be late at r, and $f_\ell^r - f_\ell^{ra\prime}$ ones could suffer from an omission at r. All symmetric faulty perceptions present in $\mathcal{V}_q(t)$ must also be present in $\mathcal{V}_r(t + \Delta t)$, however. Summing up all the differences, the expression for Δf given in Lemma 1 follows. □

3 Elementary Broadcast Primitives

In Section 4, we will provide the analysis of an advanced asynchronous algorithm, namely, the randomized Byzantine agreement algorithm of [25, 26], under our perception failure model. It relies upon two well-known broadcast primitives, *echo broadcast* [7, 26] and *simulated authenticated broadcast* [25], which have to be adapted and analyzed first. To keep the notation simple, we will use the "all-Byzantine" restriction of the perception failure model in Definition 4 ($f_s = f_o = f_c = f_m = 0$, which also implies obedient = non-faulty) throughout this section.

Echo Broadcasting

The *echo broadcast* primitive of [7, 26] implements crusader's agreement [8], which limits the power of faulty processes during a broadcast. Its interface consists of two functions *echo-broadcast*(V_p) and *echo-deliver*(V_p), which allow a process p to broadcast some value V_p to all processes in the system. Their semantics ensure that any two obedient processes that ever *echo-deliver* get the same value, and that all obedient processes will *echo-deliver* if the broadcaster is non-faulty, see Theorem 1 below.

Note that we will only consider a single instance of echo broadcasting in this section. Typical applications like the consensus algorithm of Figure 3 require multiple instances, however, which can be distinguished by their process names. We use the round number R of the application process that calls *echo-broadcast* as the process name; it is of course fixed and should not be confused with the round numbers of the echo broadcast implementation.

Figure 1 shows the pseudo code of our hybrid version of the original algorithm. It needs two concurrent single-round processes named R and R' on each processor, which send messages consisting of a type (*bcast* resp. *echo*), the process name (R resp. R'), the broadcast value V_p, and the originator of the broadcast p to each other (including itself). In the implementation of *echo-broadcast*, the broadcaster p just disseminates its value V_p to all peers. V_p will be received in process R at every processor and echoed to all. The echo messages are collected by process R', which eventually unblocks *echo-deliver* when sufficiently many have been received. The numbers f_ℓ^s, f_ℓ^r, and f_a denote the maximum

Implementation of echo-broadcast(V_p)*:*
send $(bcast, R, V_p, p)$ to all;

Process for echo-deliver(V_p)*:*
cobegin
/* Concurrent block for R */
if received $(bcast, R, V_p, p)$ from p
→ send $(echo, R', V_p, p)$ to all [once]; /* rebroadcast */
fi
/* Concurrent block for R' */
if received $(echo, R', V_p, p)$ with identical $V_p = V$ from
$n - f_\ell^s - f_\ell^r - f_a$ distinct processes
→ unblock *echo-deliver*(V_p); /* ready for *echo-deliver* at rec. */
fi
coend

Figure 1. *Echo broadcast primitive for the "all-Byzantine" hybrid failure model of Definition 4*

tolerated number of failures as specified in Definition 4. Note that f_a gives the maximum number of *processors* that may execute a faulty process R, R', or application process (calling *echo-broadcast* or *echo-deliver*). Since "at most f faulty processors" implies "at most f faulty processes (of the same name)", however, we will use the phrases "faulty processors" and "faulty processes" synonymously.

Theorem 1 (Properties Echo Broadcasting) *In a system with $n \geq 2f_\ell^s + 2f_\ell^{ra} + 2f_\ell^r + 3f_a + 1$ processors satisfying the failure model of Definition 4, where $f_a \geq 0$ gives the maximum number of Byzantine faulty processors during the entire execution, the echo broadcast primitive of Figure 1 guarantees:*

(C) Uniform Correctness*: If non-faulty processor p executes* echo-broadcast(V_p) *at time t, then* echo-deliver(V_p) *at every obedient processor is unblocked within $[t + 2\tau^-, t + 2\tau^+]$.*

(U) Uniform Unforgeability*: If processor p is obedient and does not execute* echo-broadcast(V_p) *by time t, then* echo-deliver(V_p) *cannot unblock at any obedient processor by $t + 2\tau^-$ or earlier.*

(A) Uniform Agreement*: If* echo-deliver(V_q^p) *and* echo-deliver(V_r^p) *both return a value broadcast by processor p at two obedient processes q and r, respectively, then $V_q^p = V_r^p$.*

System-wide, at most $n + 1$ broadcasts of $(\log_2 n + \log_2 C_R + \log_2 C_V + 1)$-bit messages are performed by obedient processes, where C_R resp. C_V give the cardinality of the process name space resp. the set of broadcast values.

Proof: (*Uniform Correctness*.) Since the broadcaster p is non-faulty, at least $n - f_\ell^s - f_a$ non-faulty receivers get $(bcast, R, V_p, p)$ in process R and emit $(echo, R', V_p, p)$ within time $[t + \tau^-, t + \tau^+]$, according to (A1^s) in Definition 4 and the definition of τ^-, τ^+. Consequently, any obedient receiver gets at least $n' = n - f_\ell^s - f_\ell^r - f_a$ correct $(echo, R', V_p, p)$ from different processes in process R' within another $\boldsymbol{\tau} = [\tau^-, \tau^+]$, by (A1^r) and (P1) in Definition 4. According to Figure 1, *echo-deliver*(V_p) at any obedient receiver is hence unblocked as asserted. Note that *echo-deliver*(V_p) cannot succeed before $2\tau^-$, since this could only happen if p was faulty, cp. the unforgeability proof below.

(*Uniform Unforgeability*.) The proof is by contradiction: Assume that there is an obedient process q that unblocks *echo-deliver* by $t + 2\tau^-$, which implies

$$|\mathcal{V}_q^{R'}(t+2\tau^-)| \geq n - f_\ell^s - f_\ell^r - f_a \geq f_\ell^s + 2f_\ell^{ra} + f_\ell^r + 2f_a + 1$$

according to the second **if** in Figure 1. Since only at most $f_\ell^{ra} + f_a$ of the corresponding messages $(echo, R', V_p, p)$ might originate from arbitrary receive link failures and Byzantine faulty processors, and at most f_ℓ^{sa} obedient processes could have sent $(echo, R', V_p, p)$ in response to some spurious $(bcast, R, V_p, p)$ messages caused by broadcast link failures in process R, at least one obedient processor r not affected by a broadcast link failure in process R must have sent $(echo, R', V_p, p)$, by time $t + \tau^-$. This can only happen if r got a true $(bcast, R, V_p, p)$—not a spurious one caused by an arbitrary broadcast link failure—in the first **if** sent by time t. This contradicts the assumption of the unforgeability property, however.

(*Uniform Agreement*.) If two different obedient processes q and r *echo-deliver* two different values, they must have got sufficiently many same echo-messages with different values $V_q^p \neq V_r^p$ each. We use a simple pigeonhole principle argument to show that this is impossible, given that there are only n processes that could have sent such messages: Obviously, $g_q \geq n - f_\ell^s - f_\ell^r - f_a - f_\ell^{ra} - f_a$ resp. $g_r \geq n - f_\ell^s - f_\ell^r - f_a - f_\ell^{ra} - f_a$ of the messages received at q resp. r must originate from non-faulty processes. Since there are at most $g \leq n - f_a$ such processes, $X = g_q + g_r - g$ must satisfy

$$X \geq n - 2(f_\ell^s - f_\ell^r - 2f_a - f_\ell^{ra}) + f_a \geq 1.$$

Consequently, at least one non-faulty process must have sent different messages to q and r, which is impossible.

As far as the message complexity of our algorithm is concerned, it is of course impossible to bound the number of message broadcasts by faulty processes. Every of the at most n processes that faithfully executes the algorithm of Figure 1, however, performs at most one broadcast of $(echo, R', V_p, p)$ in response to the initial one, where only process p broadcasts the message $(bcast, R, V_p, p)$. This completes the proof of Theorem 1. □

Simulated Authenticated Broadcasting

We now turn to the more involved *simulated authenticated broadcast* primitive of [25], which implements authenticated reliable broadcasts without cryptography. Note that we are dealing with the asynchronous version here; its synchronous counterpart has been analyzed in the context of the consensus algorithm of Srikanth & Toueg in [6].

Simulated authenticated broadcasting is implemented by means of two functions, namely, *sa-broadcast*(V_p) and *sa-deliver*(V_p), which allow a process p to reliably broadcast some value V_p to all processes in the system. The semantics of simulated authenticated broadcasting is captured by three properties, namely, correctness, unforgeability, and relay, defined in Theorem 2 below.

Implementation of sa-broadcast(V_p)*:*
send $(bcast, R, V_p, p)$ to all;

Process for sa-deliver(V_p)*:*
cobegin
/* Concurrent block for R */
if received $(bcast, R, V_p, p)$ from p
  → send $(echo, R', V_p, p)$ to all [once]; /* rebroadcast */
fi
/* Concurrent block for R' */
if received $(echo, R', V_p, p)$ with identical $V_p = V$ from
$f_\ell^{sa} + f_\ell^{ra} + f_a + 1$ distinct processes
  → send $(echo, R', V_p, p)$ to all [once]; /* sufficient evidence */
fi
if received $(echo, R', V_p, p)$ with identical $V_p = V$ from
$f_\ell^{sa} + 2f_\ell^{ra} + f_\ell^r + 2f_a + 1$ distinct processes
  → unblock *sa-deliver*(V_p); /* ready for *sa-deliver* at receiver */
fi
coend

Figure 2. *Simulated authenticated broadcast primitive for the "all-Byzantine" hybrid failure model of Definition 4*

Figure 2 shows the pseudo code of our hybrid simulated authenticated broadcast primitive. It consists of two concurrent single-round processes named R and R' on each processor, which send messages consisting of a type (*bcast* resp. *echo*), the process name (R resp. R'), the broadcast value V_p, and the originator of the broadcast p to each other (including itself). The initial message $(bcast, R, V_p, p)$ is used by the broadcaster p's process R to signal that the function *sa-broadcast*(V_p) has been called, $(echo, R', V_p, p)$ is

emitted (at most once) either by process R upon reception of $(bcast, R, V_p, p)$ from p, or when process R' got $(echo, R', V_p, p)$ from at least one non-faulty process ("sufficient evidence"). The figures f_a, f_ℓ^r, and f_ℓ^{ra} denote the maximum tolerated number of failures as specified in Definition 4. As in Section 3, those numbers give the maximum number of *processors* that may execute a faulty process R, R', or application process (calling *sa-broadcast* or *sa-deliver*).

Theorem 2 (Properties Simulated Auth. Broadcasting) *In a system with $n \geq f_\ell^{sa} + f_\ell^s + 2f_\ell^{ra} + 2f_\ell^r + 3f_a + 1$ processors satisfying the failure model of Definition 4, where $f_a \geq 0$ gives the maximum number of Byzantine faulty processors during the entire execution, the simulated authenticated broadcast primitive of Figure 2 guarantees:*

- *(C)* Uniform Correctness*: If non-faulty processor p calls* sa-broadcast(V_p) *at time t, then* sa-deliver(V_p) *at every obedient processor is unblocked within $[t + 2\tau^-, t + 2\tau^+]$.*
- *(U)* Uniform Unforgeability*: If processor p is obedient and does not execute* sa-broadcast(V_p) *by time t, then* sa-deliver(V_p) *cannot unblock at any obedient processor by $t + 2\tau^-$ or earlier.*
- *(R)* Uniform Relay*: If* sa-deliver(V_p) *at an obedient processor is unblocked at time t, then every obedient processor does so by time $t + \tau_\Delta$, where $\tau_\Delta = \varepsilon + \tau^+$.*

System-wide, at most $n + 1$ broadcasts of $(\log_2 n + \log_2 C_R + \log_2 C_V + 1)$-bit messages are performed by obedient processes, where C_R resp. C_V give the cardinality of the process name space resp. the set of broadcast values.

Proof: *See [21].*

It is important to note that the simulated authenticated broadcast primitive of Figure 2—unlike echo broadcasting—does *not* provide (uniform) *uniqueness* defined as: If a non-faulty (obedient) process unblocks *sa-deliver*(V_p) for a value V_p broadcast in round R, then no non-faulty (obedient) process unblocks *sa-deliver*(V_p') for a value $V_p' \neq V_p$ broadcast in round R. In the algorithm of Figure 2, a faulty broadcaster could hence "inject" multiple values in the same round, simply by inconsistently sending those to different echoing processes. Uniform relay (R) guarantees, however, that every obedient process eventually gets every value. This fact will allow us to get rid of the costly "proof-concept" in the original randomized Byzantine agreement algorithm of [26].

4 Randomized Byzantine Agreement

Enabled by the results of the previous section, we are ready for investigating the randomized consensus algorithm of [25] under our perception failure model. A (randomized) consensus—also called Byzantine agreement—algorithm computes (with high probability) a common decision value V based upon initial values v_p provided locally at every process p. Our algorithm will be based upon Toueg's improvement [26] of the algorithm proposed by Rabin in [19], which uses authenticated broadcasts and Shamir's secret sharing scheme [24]. By plugging in the hybrid broadcast primitives from Section 3, a hybrid version of this algorithm is easily derived.

The secret sharing scheme of [24] assumes a non-faulty *dealer* D, which generates a sequence of random bits $s_1, s_2, \ldots$ and, for each bit s_k, n *pieces* s_k^i, $1 \leq i \leq n$. Those pieces are such that the knowledge of $t + 1$ of those is necessary and sufficient for computing s_k. The dealer signs all pieces with its signature σ_D to prevent forgery and distributes to each process i the sequence of its pieces $\sigma_D(s_1^i), \sigma_D(s_2^i), \ldots$. Note that this happens off-line, *prior* to the execution(s) of the actual consensus algorithm, and is the only place where authentication will be required in our hybrid algorithm. This secret sharing scheme makes it impossible for t faulty processes to compute the secret solely from their pieces at runtime — at least one non-faulty process's piece must be available for this purpose as well.

The original algorithm of [26] (cp. Figure 3) computes, with probability at least $1 - (1/2)^K$, a common consensus value V from all the processes' input values $v_p \in \{0, 1\}$, $1 \leq p \leq n$, by means of K iterations with three phases (which may consist of multiple rounds each). In phase 1, all processes broadcast their current consensus value V_p (initially $V_p = v_p$) and wait for the arrival of $n - f$ authenticated messages from different processes (including itself). Every process then computes a new value for V_p based upon the values in the received messages and saves the latter as a *proof* for its choice.

In phase 2, every process *echo-broadcasts* its new value V_p along with its proof to all processes in the system and waits for $n - f$ such messages from different processes. It computes the number of messages containing $V_q = 1$ among those and saves them in the variable *count*.

In phase 3, every process p in iteration k discloses its piece of the secret s_k^p by broadcasting it to all processes, and waits for the arrival of $f + 1$ pieces—with a correct dealer's signature—from different processes. When they arrive, process p can compute the shared secret s_k. Note that it does not matter whether a piece comes from a faulty or non-faulty process since forging a piece is impossible due to authentication.

Finally, a new consensus value V_p is computed by suitably combining the random bit s_k with the value of *count* obtained in phase 2: The algorithm ensures that, with probability at least 1/2, all non-faulty processes achieve the same new consensus value at the end of an iteration, even if there

was arbitrary disagreement before.

The Hybrid Algorithm

We will now develop a hybrid variant of the above algorithm for the "all-Byzantine" setting of our perception failure model. It is assumed here that there are $n \geq 2f_\ell^s + 2f_\ell^{ra} + 2f_\ell^r + 3f_a + 1$ processors in the system, with at most $f_a \geq 0$ arbitrary faulty ones among those. Note that the "all-Byzantine" setting is appropriate here, since crash failures are as severe as arbitrary failures for our algorithm.

The pseudo code of our hybrid randomized consensus algorithm is shown in Figure 3. Note that link failures show up explicitly only in phase 3 of the hybrid algorithm, since they are completely hidden by the broadcast primitives in phases 1 and 2.

Code for process p:
$V_p = v_p$; /* Initial value */
for $k = 0$ **to** $K - 1$ **do**

/* Phase 1 (round $R = 3k$) */
sa-broadcast(V_p);
wait for *sa-deliver*(V_q) from $n - f_a$ processes;
count := number of q's with $V_q = 1$;
if *count* $\geq n - 2f_a$ **then** $V_p := 1$ **else** $V_p := 0$;
fi

/* Phase 2 (round $3k + 1$) */
echo-broadcast(V_p);
wait for *echo-deliver*(V_q) with *acceptable*(V_q) from $n - f_a$ procs;
count := number of q's with $V_q = 1$;

/* Phase 3 (round $3k + 2$) */
send$(\sigma_D(s_k^p))$ to all; /* disclose piece */
wait for *receive*$(\sigma_D(s_k^q))$ with correct σ_D from $f_\ell^{ra} + f_a + 1$ distinct processes;
compute s_k from the received pieces;
if ($s_k = 0$ **and** count ≥ 1) **or**
($s_k = 1$ **and** count $\geq 2f_a + 1$) **then** $V_p = 1$;
else $V_p := 0$;
fi
od

Figure 3. *Randomized binary consensus algorithm for the hybrid failure model of Definition 4*

Figure 3 reveals that we only replaced the authenticated broadcast in phase 1 of the original algorithm by the simulated broadcast primitive of Section 3. The latter does not employ signatures, however, so we could not retain the *proof* concept. Instead, we exploit the fact that simulated authenticated broadcasting satisfies the relay property (R) according to Theorem 2: It guarantees that, eventually, every non-faulty process must get all messages seen by any other non-faulty process (although not necessarily in the same sequence). As as consequence, our algorithm needs to *echo-broadcast* the single value V_p only, which considerably reduces the communication costs.

Receiver q can verify in phase 2 whether the value V_p disseminated by some process p via *echo-broadcast* is legitimate as follows: It just looks whether there are $n - f_a$ phase 1 messages among its—possibly larger—set of received ones, such that V_p satisfies the condition of the **if** of phase 1 if applied to those; we express this via the predicate *acceptable*(V_p). It is important to note, however, that q must wait until $n - f_a$ phase 2 messages from different p's have passed this test — after all, it may be the case that the phase 1 messages that were used by p to compute its V_p in phase 2 did not yet arrive at q.

By means of a proof that almost literally follows the original one in [26], it is not difficult to establish the following major Theorem 3:

Theorem 3 (Properties Randomized Consensus) *In a system with $n \geq 2f_\ell^s + 2f_\ell^{ra} + 2f_\ell^r + 3f_a + 1$ processors according to Definition 4, where $f_a \geq 0$ denotes the maximum number of arbitrary faulty processors during the whole execution, the randomized consensus algorithm of Figure 3 satisfies:*

P1. Termination*: All non-faulty processes terminate the algorithm.*

P2. Validity*: If all non-faulty processes p start with the initial value $v_p = m$, then every non-faulty process terminates the algorithm with $V = m$.*

P3. Randomized Agreement*: With probability at least $1 - (1/2)^K$, every non-faulty process terminates the algorithm with the same value V.*

Proof: *See [21].*

5 Stubborn Channels

The analysis in Section 4 showed that our algorithm tolerates a considerable number of link faults without additional measures, provided that sufficiently many processors are available. For example, pure UDP datagram communication can be used instead of TCP in systems with moderate link failure rates. For high link failure rates, it is quite likely that our link failure bounds f_ℓ^s and f_ℓ^r are too restrictive [21]. Resorting to some kind of reliable communication might be inevitable. However, perfect communication requires unbounded memory space [10].

Stubborn channels [14] have been proposed as an alternative. A *k-stubborn channel* is a point-to-point communication link that reliably delivers the last k messages submitted to it for transmission, provided that both sender

and receiver are non-faulty and the sender eventually stops submitting messages (such that "last k messages" makes sense). Obviously, a 1-stubborn channel can easily be implemented atop of datagrams by using a single buffer: The message in the buffer is periodically retransmitted until an acknowledgment is received. If a new message is submitted for transmission before the previous one has been acknowledged, it just overwrites the previous message in the buffer. A k-stubborn channel can be implemented by using k 1-stubborn channels operating on a circular buffer.

In spite of being powerful enough for solving consensus in asynchronous systems [14], a k-stubborn channel needs only bounded memory space. Our algorithm reconfirms this fact, since it is easily modified to work with 2-stubborn channels. The modification required is *forcing* the execution of the algorithm to some future iteration: A process currently in iteration k is forced to iteration $k' > k$ by interrupting the current execution, setting the iteration loop counter to k', and resuming execution in the wait of phase 1 in iteration k'. Note, however, that *sa-broadcast*$(V_p^{k'})$ is not called when forcing process p.

The following Lemma 2 shows that forcing does not affect the consensus result, provided that at least one non-faulty process is ahead by two iterations:

Lemma 2 (Forcing Rounds) *Suppose some non-faulty process p_l is currently in iteration $k' \leq k-2$ at time t, when some other non-faulty process p_a is already in iteration $k \geq 2$. Then, all non-faulty processes will complete iteration $k-1$ and thus enter iteration k also when p_l is forced to iteration $k-1$. The consensus value resp. the probability of reaching agreement is not affected by forcing process p_l.*

Proof: *See [21].*

The result of Lemma 2 implies that process p_l can be forced to iteration $k-1$ at time t, without changing the outcome of iteration $k-1$, when sufficient evidence for the existence of a non-faulty process in iteration k is obtained. This is the case when iteration k messages from $f_a + f_\ell^{ra} + 1$ distinct processes arrive, since only at most f_ℓ^{ra} of those could be spurious messages from arbitrary receive link faults, and f_a messages could originate from faulty processes.

Most importantly, if round forcing is employed, there is no need for a process in iteration k to support iteration k' for any $k' \leq k-2$. In particular, all echo broadcasting and simulated authenticated broadcasting processes belonging to iteration k' can be killed upon switching to round k. After all, it is guaranteed that all non-faulty processes will eventually complete iteration $k-1$ and enter iteration k, which means that all late processes—which might not terminate since some forced processes stopped their support of earlier rounds—must eventually get sufficient evidence and be forced to iteration $k-1$.

Since a process in iteration k needs to deal with messages belonging to iteration k or $k-1$ only, it is not difficult to show that it suffices to reliably transmit (and receive) only the two highest-round messages of a given type if forcing is employed, i.e., that 2-stubborn channels will be sufficient. This will be done in Theorem 4 below, where we assume that each processor executes $2n$ "generic" processes [*sa-R*] resp. [*sa-R'*], which implement all the iteration's (dedicated) processes R resp. R' of simulated authenticated broadcasting for the n peer processors; recall that one dedicated process per iteration and processor was assumed in Section 4. Similarly, we need $2n$ generic processes [*echo-R*] resp. [*echo-R'*] on each processor for echo broadcasting. Finally, one process per processor executes the randomized Byzantine agreement algorithm. Note that all those processes are always in the same round, and are all forced together.

Theorem 4 (Stubborn Channels) *Theorem 3 remains valid if every pair of instances of the algorithms of Figure 3, 2 and 1 employs a dedicated 2-stubborn channel for basic communication, provided that round forcing is applied when (1) $f_a + f_\ell^{ra} + 1$* sa-broadcast *or* echo-broadcast *init-messages* $(\text{bcast}, R, V_p, p)$ *arrived from distinct processors, or (2) $(f_a + f_\ell^{ra} + 1) \times (f_a + f_\ell^{ra} + 1)$ distinct echo-messages* $(\text{echo}, R', V_p, p)$ *arrived, as witnesses of $f_a + f_\ell^{ra} + 1$ distinct* $(\text{bcast}, R, V_p, p)$.

Proof: *See [21].*

Note that (the proof of) Theorem 4 also implies that any process q needs to store only the two perceptions $V_q^{p,R}$ with the highest round numbers received from any p, i.e., one does not need unbounded memory space for perception vectors.

6 Conclusions

We introduced a novel hybrid failure model for round-based distributed algorithms in partially synchronous systems with possibly unknown delays. It accommodates both process and link failures and distinguishes asymmetric, symmetric, omission, clean crash, and manifest failures, both in the time and in the value domain. Our perception-based failure model considerably simplifies accurate fault-tolerance analysis, yet allows the evaluation of assumption coverage and running times as well [21]. It is hence well-suited for both synchronous and asynchronous wireline and, in particular, wireless networked systems.

We analyzed a hybrid version of the randomized Byzantine agreement of Srikanth & Toueg, which is based upon suitably adopted variants of the asynchronous echo broadcast and simulated authenticated broadcast primitives. Its

probability of disagreement was found to be only 2^{-K}, which is the same as for asynchronous systems without link failures. With respect to link failure tolerance, it turned out that resource redundancy (more processes) can be used instead of time redundancy (retransmissions) for this purpose: Tolerating f_ℓ^r resp. f_ℓ^s receive resp. send link failures at every node, in every round, with $f_\ell^{ra} \leq f_\ell^r$ arbitrary ones among the f_ℓ^r receive link failures, just needs $2f_\ell^s + 2f_\ell^r + 2f_\ell^{ra}$ additional nodes. Although a comparison with the lower bound $f_\ell^s + f_\ell^{sa} + f_\ell^r + f_\ell^{ra}$ from [22] reveals that this is sub-optimal (whereas the resilience with respect to process failures is optimal [15]), our algorithm can nevertheless cope with up to $\mathcal{O}\big(Kn^2\big)$ link failures systemwide during K iterations. For excessive link failure rates, 2-stubborn links, which avoid unbounded memory space required for implementing perfect communications, can also be used.

References

[1] Y. Afek, H. Attiya, A. Fekete, M. Fischer, N. Lynch, Y. Mansour, D.-W. Wang, and L. Zuck. Reliable communication over unreliable channels. *Journal of the ACM (JACM)*, 41(6):1267–1297, 1994.

[2] M. K. Aguilera, W. Chen, and S. Toueg. On quiescent reliable communication. *SIAM Journal of Computing*, 29(6):2040–2073, April 2000.

[3] M. Azadmanesh and R. M. Kieckhafer. Exploiting omissive faults in synchronous approximate agreement. *IEEE Transactions on Computers*, 49(10):1031–1042, Oct. 2000.

[4] K. Bartlett, R.A.Scantlebury, and P. Wilkinson. A note on reliable full-duplex transmission over half-duplex links. *Communications of the ACM (JACM)*, 12(5):260–261, 1969.

[5] A. Basu, B. Charron-Bost, and S. Toueg. Crash failures vs. crash + link failures. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, page 246. ACM Press, 1996.

[6] M. Biely and U. Schmid. Message-efficient consensus in presence of hybrid node and link faults. Technical Report 183/1-116, Department of Automation, Technische Universität Wien, August 2001. (submitted).

[7] G. Bracha and S. Toueg. Resilient consensus protocols. In *Proceedings of the 2nd Symposium on the Principles of Distributed Computing (PODC'83)*, pages 12–26, Montreal, Canada, 1983.

[8] D. Dolev. The Byzantine generals strike again. *Journal of Algorithms*, 3(1):14–30, 1982.

[9] C. Dwork, N. Lynch, and L. Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM*, 35(2):288–323, Apr. 1988.

[10] A. Fekete, N. Lynch, Y. Mansour, and J. Spinelli. The impossibility of implementing reliable communication in the face of crashes. *Journal of the ACM (JACM)*, 40(5):1087–1107, 1993.

[11] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty processor. *Journal of the ACM*, 32(2):374–382, Apr. 1985.

[12] E. Gafni. Round-by-round fault detectors (extended abstract): unifying synchrony and asynchrony. In *Proceedings of the Seventeenth Annual ACM Symposium on Principles of Distributed Computing*, pages 143–152. ACM Press, 1998.

[13] J. N. Gray. Notes on data base operating systems. In G. S. R. Bayer, R.M. Graham, editor, *Operating Systems: An Advanced Course*, volume 60 of *Lecture Notes in Computer Science*, chapter 3.F, page 465. Springer, New York, 1978.

[14] R. Guerraoui, R. Oliveira, and A. Schiper. Stubborn communication channels. Technical Report 98-278, Dèpartment d'Informatique, Ècole Polytechnique Fédérale de Lausanne, 1998.

[15] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.

[16] G. Le Lann. Certifiable critical complex computing systems. In K. Duncan and K. Krueger, editors, *Proceedings 13th IFIP World Computer Congress 94*, volume 3, pages 287–294. Elsevier Science B.V. (North-Holland), 1994.

[17] G. Le Lann and U. Schmid. How to implement a timer-free perfect failure detector in partially synchronous systems. Technical Report 183/1-127, Department of Automation, Technische Universität Wien, January 2003. (submitted).

[18] N. Lynch. *Distributed Algorithms*. Morgan Kaufman, 1996.

[19] M. O. Rabin. Randomized Byzantine Generals. In *Proceedings of the 24th Annual IEEE Symposium on Foundations of Computer Science*, pages 403–409, 1983.

[20] U. Schmid. How to model link failures: A perception-based fault model. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN'01)*, pages 57–66, Göteborg, Sweden, July 1–4, 2001.

[21] U. Schmid and C. Fetzer. Randomized asynchronous consensus with imperfect communications. Technical Report 183/1-120, Department of Automation, Technische Universität Wien, January 2002.

[22] U. Schmid and B. Weiss. Synchronous Byzantine agreement under hybrid process and link failures. Technical Report 183/1-124, Department of Automation, Technische Universität Wien, Nov. 2002. (submitted; replaces TR 183/1-110).

[23] U. Schmid, B. Weiss, and J. Rushby. Formally verified byzantine agreement in presence of link faults. In *22nd International Conference on Distributed Computing Systems (ICDCS'02)*, pages 608–616, July 2-5, 2002.

[24] A. Shamir. How to share a secret. *Communications of the ACM*, 22:612–613, 1979.

[25] T. Srikanth and S. Toueg. Simulating authenticated broadcasts to derive simple fault-tolerant algorithms. *Distributed Computing*, 2:80–94, 1987.

[26] S. Toueg. Randomized byzantine agreements. In *Proceedings of the 3rd Symposium on Principles of Distributed Computing (PODC'84)*, pages 163–178, 1984.

[27] C. J. Walter and N. Suri. The customizable fault/error model for dependable distributed systems. *Theoretical Computer Science*, 290:1223–1251, October 2002.