

Communication-Optimal Failure Detection in Partially Synchronous Systems [☆]

Alberto Lafuente^{*,a}, Mikel Larrea^a, Iratxe Soraluze^a, Roberto Cortiñas^a

^a*University of the Basque Country, Faculty of Computer Science,
Lardizabal 1, 20018 San Sebastián, Spain*

Abstract

Since Chandra and Toueg proposed the failure detector abstraction for crash-prone systems, several algorithms implementing failure detectors in partially synchronous systems have been proposed. *Communication efficiency* is a performance measure used to evaluate failure detection algorithms, and it is defined as the number of links carrying messages forever. In this regard, communication-efficient $\diamond\mathcal{P}$ algorithms in which just n links carry messages forever have been proposed, being n the number of processes in the system. These algorithms are based on heartbeat messages circulating in a ring arrangement of processes. In this paper, we present *communication optimality*, a communication efficiency degree reached when only \mathcal{C} links carry messages forever, being \mathcal{C} the number of correct processes. We show that \mathcal{C} is the minimum number of links carrying messages forever required to implement

[☆]Research partially supported by the Spanish Research Council, under grants TIN2007-67353-C02-02 and TIN2006-15617-C03-01, the Basque Government, under grant S-PE07IK03, and the Comunidad de Madrid, under grant S-0505/TIC/0285.

*Corresponding author. Telephone: +34 943015086. Fax: +34 943015590.

Email addresses: `alberto.lafuente@ehu.es` (Alberto Lafuente),
`mikel.larrea@ehu.es` (Mikel Larrea), `iratxe.soraluze@ehu.es` (Iratxe Soraluze),
`roberto.cortinas@ehu.es` (Roberto Cortiñas)

$\diamond\mathcal{P}$ and that communication-optimal $\diamond\mathcal{P}$ algorithms can be implemented. We also show that, assuming that there is at least one incorrect process, \mathcal{C} is optimal not only for $\diamond\mathcal{P}$ but also for $\diamond\mathcal{S}$ and Ω . Finally, we propose two communication-optimal $\diamond\mathcal{P}$ algorithms that follow two different approaches to communicate suspicions. One of them uses a one-to-all communication pattern, while the other uses exclusively one-to-one communication. A trade-off results between detection latency and sporadic communication overhead. Depending on the characteristics of the system, the most suitable approach can be followed.

Key words: distributed algorithms, fault tolerance, consensus, partial synchrony, unreliable failure detectors, communication optimality

Footnotes:

(1) In the algorithm in [21], for every process (correct or faulty), the link coming from its correct predecessor in the ring is used permanently.

(2) We assume that messages include the identity of the sender and a sequence number, which make every message unique.

(3) An optimization consists in not relaying m to p , $sender(m)$ and the process q from which m has been received for the first time (if $q \neq sender(m)$).

(4) By equivalence, the reasoning regarding Ω applies to $\diamond\mathcal{S}$ too.

(5) Here we informally use the terms *nearest predecessor* (or *nearest successor*) of a process p to denote the first process preceding (or succeeding) p following the ring arrangement and fitting a particular condition.

1. Introduction

1.1. Background and Related Work

Unreliable failure detectors, proposed by Chandra and Toueg in [1], are a mechanism providing (possibly incorrect) information about process failures. This mechanism has been used to solve several problems in asynchronous crash-prone distributed systems, in particular the *consensus* problem [2]. A class of failure detectors is defined by a *completeness* property, which characterizes the failure detector's capability of suspecting incorrect processes, and by an *accuracy* property, which restricts the mistakes the failure detector can make. More specifically, Chandra and Toueg defined, among others, the following two completeness properties and two accuracy properties that a failure detector may satisfy:

- *Strong Completeness*: eventually every process that crashes is permanently suspected by *every* correct process.
- *Weak Completeness*: eventually every process that crashes is permanently suspected by *some* correct process.
- *Eventual Strong Accuracy*: there is a time after which correct processes are not suspected by any correct process.
- *Eventual Weak Accuracy*: there is a time after which some correct process is never suspected by any correct process.

The combination of these properties gives us four classes of failure detectors, which are shown in Table 1. Consensus can be solved using a failure

	Eventual Strong Accuracy	Eventual Weak Accuracy
Strong Completeness	<i>Eventually Perfect</i> ($\diamond\mathcal{P}$)	<i>Eventually Strong</i> ($\diamond\mathcal{S}$)
Weak Completeness	<i>Eventually Quasi Perfect</i> ($\diamond\mathcal{Q}$)	<i>Eventually Weak</i> ($\diamond\mathcal{W}$)

Table 1: Four classes of failure detectors.

detector of any of those four classes. In particular, there is another failure detector class denoted Ω , that is equivalent to $\diamond\mathcal{S}$ and $\diamond\mathcal{W}$, and has been proved to be the weakest for solving consensus [3]. The Ω failure detector class provides eventual agreement on a common and correct *leader* among all non-faulty processes in a system.

Specific algorithms for implementing Ω and/or $\diamond\mathcal{S}$ have been proposed, e.g. [4–15]. Observe that $\diamond\mathcal{P}$ trivially satisfies the properties of $\diamond\mathcal{S}$. Also, $\diamond\mathcal{P}$ can be easily transformed into Ω , e.g., by choosing as leader the non-suspected process with the lowest identifier.

Failure detectors of the class $\diamond\mathcal{P}$, being strictly stronger than Ω and $\diamond\mathcal{S}$, can also be used to solve consensus, providing better accuracy. Several algorithms implementing $\diamond\mathcal{P}$ have been proposed in the literature. The algorithm proposed by Chandra and Toueg in [1] uses a heartbeat mechanism and all-to-all communication to detect faulty processes. The algorithms proposed by Aguilera et al. in [4] and by Larrea et al. in [16] use heartbeats too, and rely on a leader-based approach. On the other hand, the algorithms proposed by Larrea et al. in [17, 18] use a polling —or query/reply— mechanism on a ring arrangement of processes. Roughly speaking, the leader-based and the ring-based algorithms are more efficient than the all-to-all algorithm regarding the number of messages sent (linear vs. quadratic). Observe also that, compared to polling, the heartbeat mechanism reduces the number of

messages to the half. Moreover, heartbeats inherently provide a certain level of communication reliability in a system with fair lossy links, while polling usually requires reliable communication.

Chandra and Toueg showed in [1] that $\diamond Q$ can also be used to solve consensus. To do so, they first showed that classes $\diamond Q$ and $\diamond P$ are equivalent from a problem solvability point of view, i.e., a problem which is solvable with $\diamond P$ is also solvable with $\diamond Q$ and vice versa. It is worth noting that the equivalence of $\diamond Q$ and $\diamond P$ does not come for free, i.e., not all failure detectors in $\diamond Q$ are in $\diamond P$. Instead, it means that any failure detector in $\diamond Q$ can be extended with a simple distributed algorithm to obtain a failure detector in $\diamond P$.

Since consensus cannot be solved in crash-prone, pure asynchronous systems [19], algorithms that implement failure detectors make some weak timing assumptions. Specifically, a partially synchronous model [1, 20] is considered in this work. In such a model, in every run of the system, there are bounds on relative process speeds and on message transmission times, but these bounds are not known and they hold only after some unknown but finite time. In practice, the bounds depend on parameters such as the network speed or the size of the system, and hold easier in, for example, local area networks than in wide area networks. Actually, the bounds must exist and hold only for the links that connect correct processes. Hence, it is important to design failure detection algorithms that use a low number of links, e.g., by arranging the processes in a logical ring topology.

As shown recently in [21], algorithms for the class $\diamond P$ combining a heartbeat-based detection mechanism on a logical ring arrangement of pro-

cesses outperform the aforementioned ones in terms of the number of links permanently used, while preserving good quality-of-service. In this regard, algorithms in [21] are communication-efficient following Aguilera et al. [4], i.e., eventually only n unidirectional links carry messages forever. With regard to this performance measure, heartbeat-based ring algorithms outperform other ring algorithms based on polling [17, 18] or algorithms using a centralized communication pattern [4, 16]. By all means, algorithms using an all-to-all communication pattern, such as Chandra-Toueg’s algorithm [1], are far from being communication-efficient.

Other failure detection algorithms specifically designed for wide area networks, e.g. [22], are based on local failure detectors that provide their properties in a neighborhood of processes (using all-to-all communication inside each neighborhood), and propagate the information about suspicions among neighborhoods. The advantage of a ring based approach is that, once the logical ring is defined, neighborhoods are implicit, and eventually involve just two processes for every correct process, i.e., its correct predecessor and correct successor in the ring.

1.2. Implementing Failure Detectors

Obviously, an algorithm implementing a failure detector of a given class must satisfy the properties defined for that class. When implementing a failure detector, besides satisfying the properties of the class it belongs to, performance should be also taken into account. In this work, we focus on the following two performance issues:

- i) Scalability, to allow a failure detector to be deployed in networks with a high number of nodes and/or heterogeneous links. A scalable failure

detection algorithm should use a low number of links and avoid all-to-all communication patterns.

- ii) Quality-of-service, involving several parameters, such as detection latency or stabilization time [23], which affect the responsiveness of the system. For example, when a crashed process q is suspected by a process p , every correct process should be informed as soon as possible in order to provide low detection latencies. In this regard, a one-to-all communication pattern for suspicion propagation can be helpful. However, when the suspicion is erroneous such a communication pattern could become a drawback, since it propagates the erroneous suspicion in the system.

Keeping in mind scalability and quality-of-service, we define now a set of interesting properties for an algorithm that implements a failure detector:

Communication efficiency. In a communication-efficient algorithm only n unidirectional links carry messages forever, being n the number of processes in the system [4]. Observe that communication efficiency refers to a *permanent* behavior that will hold eventually. In large systems, communication-efficient algorithms will be more scalable than non communication-efficient algorithms.

Low sporadic overhead. Besides the permanent communication cost due to periodic messages (also known as heartbeats), which is addressed by the previous communication efficiency property, an algorithm implementing a failure detector can involve *sporadic* extra messages caused

by failure suspicions, resulting in a peak overhead. There is a trade-off in the way this sporadic traffic is managed. On the one hand, a one-to-one communication pattern can be used (e.g., by using a logical ring arrangement); this way communication overhead is reduced and, henceforth, it provides better scalability. On the other hand, one-to-all (or even all-to-all) communication leads to provide low detection latencies.

Communication locality. This property is two-folded. On the one hand, we will say that a failure detection algorithm has *periodic communication locality* when periodic, permanent monitoring messages (i.e., heartbeats), are sent to some process(es) in the neighborhood of the sender process. Observe that a logical ring arrangement of processes based on proximity naturally provides this property. Similarly, we will say that a failure detection algorithm has *sporadic communication locality* when messages sent as a consequence of a suspicion are sent likely to some process(es) in the neighborhood of the suspecting process. Again, trade-offs should be considered between the good scalability associated to communication locality, and the potential good quality-of-service provided by the use of one-to-all and all-to-all communication patterns.

1.3. Using Communication-Optimal $\diamond\mathcal{P}$ to Solve Consensus

Despite $\diamond\mathcal{S}$ and Ω have been extensively used to solve consensus [1, 24–27], we focus our work on implementing *communication-optimal* failure detectors of the class $\diamond\mathcal{P}$. As we will show, in a communication-optimal $\diamond\mathcal{P}$ algorithm only \mathcal{C} unidirectional links carry messages forever, being \mathcal{C} the

number of correct processes. The choice of $\diamond\mathcal{P}$ is mainly justified by the fact that, as we will also show, communication optimality is almost the same for $\diamond\mathcal{P}$, $\diamond\mathcal{S}$ and Ω , and a failure detector of the class $\diamond\mathcal{P}$ trivially implements $\diamond\mathcal{S}$ and Ω . Hence, consensus algorithms based on either $\diamond\mathcal{S}$, Ω , or an equivalent leader election mechanism, e.g. [16, 28–30], can also benefit of communication-optimal implementations of $\diamond\mathcal{P}$. Moreover, for certain problems [31] and consensus protocols [32] failure detector $\diamond\mathcal{P}$, being stronger than $\diamond\mathcal{S}$ and Ω , is required. Finally, failure detectors of the class $\diamond\mathcal{P}$ are more natural, in the sense that all the correct processes can provide a precise set composed of exclusively crashed processes, providing a better degree of accuracy.




Figure 1: Links used permanently in (a) a communication-efficient algorithm, and (b) a communication-optimal algorithm.

Figure 1 shows an example of the number of unidirectional links used permanently for a system composed of eight processes, out of which five are correct, i.e., $n = 8$ and $\mathcal{C} = 5$. Faulty processes are represented by grey circles. Observe that in a communication-efficient algorithm, e.g. [21], n links are used permanently,¹ while in a communication-optimal algorithm only \mathcal{C}

¹In the algorithm in [21], for every process (correct or faulty), the link coming from its

links are used permanently, which is optimal for implementing $\diamond\mathcal{P}$.

1.4. Summary of Contributions

The contributions of this paper are the following:

1. We study the communication efficiency of implementing $\diamond\mathcal{P}$ in partially synchronous systems where processes can fail by crashing, and introduce the notion of *communication optimality*.
2. We show that the minimum number of unidirectional links carrying messages forever needed for an algorithm to provide the properties of $\diamond\mathcal{P}$ is \mathcal{C} , i.e., the number of correct processes in the system.
3. We show that \mathcal{C} is also minimal for Ω when at least one process crashes. Hence, communication optimality is almost the same for $\diamond\mathcal{P}$, $\diamond\mathcal{S}$ and Ω , which makes communication-optimal $\diamond\mathcal{P}$ algorithms good candidates to be used in a consensus algorithm.
4. We show that communication-optimal $\diamond\mathcal{P}$ algorithms can be implemented. More precisely, we propose two ring-based communication-optimal $\diamond\mathcal{P}$ algorithms, one of them using a one-to-all communication pattern for communicating suspicions, the other using exclusively a one-to-one communication pattern.

With respect to the communication efficiency property, the two approaches proposed in this paper lead to communication-optimal $\diamond\mathcal{P}$ algorithms, i.e.,

correct predecessor in the ring is used permanently.

eventually only \mathcal{C} links carry messages forever. Concerning the sporadic overhead involved by a suspicion, the first algorithm has a higher overhead due to the use of a reliable broadcast communication primitive, while the second algorithm has a low overhead. Finally, regarding communication locality, the first algorithm has only periodic communication locality, while the second algorithm has both periodic and sporadic communication locality. In this regard, our notion of neighborhood is dynamic and related to the estimation of the correct predecessor and successor of a process in the ring. Note that according to the properties of $\diamond\mathcal{P}$, the neighborhood of every correct process will eventually stabilize. If the logical ring is arranged using proximity criteria, e.g., the number of physical hops between consecutive processes in a wide area network, our second algorithm minimizes the network traffic, which could provide benefits in geographically dispersed networks.

1.5. Roadmap

The rest of the paper is organized as follows. In Section 2, we describe the system model considered in this work. In Section 3, we show the communication optimality results for $\diamond\mathcal{P}$ and Ω . In Section 4, we give two communication-optimal $\diamond\mathcal{P}$ algorithms. In Section 5, we analyze the performance of the algorithms, and compare them with previously proposed $\diamond\mathcal{P}$ algorithms. Finally, Section 6 concludes the paper.

2. System Model

We consider a distributed system composed of a finite set Π of $n > 1$ processes, $\Pi = \{p_1, p_2, \dots, p_n\}$, that communicate only by sending and receiving messages. We also use the alternative notation p, q, r, \dots to denote

processes. Every pair of processes (p, q) is connected by two unidirectional and reliable logical communication links $p \rightarrow q$ and $q \rightarrow p$. This means that process p can send a message to process q using a *send* primitive and vice versa. The definition of reliable link that we consider is the following: if both the sender and the receiver do not crash, then all messages that are sent are eventually received. Reliable communication is usually implemented using retransmission techniques and acknowledgment messages.

Processes can only fail by crashing, that is, by prematurely halting. Moreover, crashes are permanent, i.e., crashed processes do not recover. In every run of the system we identify two complementary subsets of Π : the subset of processes that do not fail, denoted *correct*, and the subset of processes that do fail, denoted *crashed*. We use \mathcal{C} to denote the number of correct processes in the system in the run of interest, which we assume is at least one, i.e., $\mathcal{C} = |\text{correct}| \geq 1$.

We consider that processes are arranged in a logical ring. Without loss of generality, process p_i is preceded by process p_{i-1} , and followed by process p_{i+1} . As usual, p_1 follows p_n in the ring. In general, we will use the functions $\text{pred}(p)$ and $\text{succ}(p)$ respectively to denote the predecessor and the successor of a process p in the ring.

Concerning timing assumptions, we consider a partially synchronous model [1, 20] which stipulates that, in every run of the system, there are bounds on relative process speeds and on message transmission times, but these bounds are not known and they hold only after some unknown but finite time (called *GST* for *Global Stabilization Time*). The communication links supporting this behaviour are also called *eventually timely* links [5]. Our model is ac-

tually a variant of the partial synchrony models of [1, 20]. The difference is that we assume reliable communication links. Nevertheless, when presenting the algorithms we will refine the minimal assumptions on communication reliability and synchrony required by each algorithm.

Finally, in the algorithms presented in this paper we assume that a local clock that can measure real-time intervals is available to each process. Clocks are not synchronized.

Algorithm 1: Reliable Broadcast by message diffusion.

{ Every process p executes the following }

To execute **R-broadcast**(m):

└ send m to all (including p)

R-deliver(m) occurs as follows:

┌ **when** receive m for the first time **do**
└ ┌ **if** $\text{sender}(m) \neq p$ **then**
└ └ ┌ send m to all
└ └ └ R-deliver(m)

2.1. Reliable Broadcast

Reliable Broadcast is a communication primitive for asynchronous systems that we use in one of our algorithms. It guarantees that all correct processes deliver the same set of messages. This set includes at least all messages broadcast by correct processes. Formally, Reliable Broadcast is defined in terms of two primitives, *R-broadcast*(m) and *R-deliver*(m), and satisfies the following properties [33]:

- Validity. If a correct process R-broadcasts a message m , then it eventually R-delivers m .

- Agreement. If a correct process R-delivers a message m , then all correct processes eventually R-deliver m .
- Uniform integrity. For any message m , every process R-delivers m at most once, and only if m was previously R-broadcast by $sender(m)$.²

Algorithm 1 presents a simple Reliable Broadcast algorithm for asynchronous systems with up to $n - 1$ crash failures [1]. Informally, when a process receives a message for the first time, it relays the message to all processes and then R-delivers it.³

3. On Communication Optimality

In this section, we show that \mathcal{C} , i.e., the number of correct processes in the system, is the minimum number of unidirectional links carrying messages forever needed for an algorithm to provide the properties of $\diamond\mathcal{P}$. Then, we show that, assuming that at least one process crashes, i.e., $\mathcal{C} < n$, \mathcal{C} is also minimal for implementing Ω .⁴ Both results hold when there are at least two correct processes in the system, i.e., $\mathcal{C} \geq 2$.

Theorem 1. *\mathcal{C} is the minimum number of unidirectional links carrying messages forever needed for an algorithm to provide the properties of $\diamond\mathcal{P}$ in a crash-prone system.*

²We assume that messages include the identity of the sender and a sequence number, which make every message unique.

³An optimization consists in not relaying m to p , $sender(m)$ and the process q from which m has been received for the first time (if $q \neq sender(m)$).

⁴By equivalence, the reasoning regarding Ω applies to $\diamond\mathcal{S}$ too.

PROOF. Given a run R , observe that every process must periodically inform that it is still alive by sending a message, which after every incorrect process has crashed gives us the minimum number of \mathcal{C} unidirectional links. Otherwise, if less than \mathcal{C} unidirectional links carry messages forever, there is some correct process p that eventually stops sending messages. Let t be the time instant in which p stops sending messages. Consider now another run R' , identical to R until time t , and assume that p crashes at time t in R' . For any correct process q , if q does not eventually and permanently suspect p , then the strong completeness property of $\diamond\mathcal{P}$ is violated. Hence, q will eventually and permanently suspect p in R' . Observe that both executions R and R' are indistinguishable. Hence, in run R q will also eventually and permanently suspect p , violating the eventual strong accuracy property of $\diamond\mathcal{P}$.

Theorem 2. *If at least one process crashes, then \mathcal{C} is the minimum number of unidirectional links carrying messages forever needed for an algorithm to provide the property of Ω in a crash-prone system.*

PROOF. The proof is by contradiction. Assume that we have an implementation of Ω in which only $\mathcal{C} - 1$ unidirectional links carry messages forever. Observe that such an implementation will be possible only if correct processes are arranged in a tree topology, being the leader the root of the tree and propagating heartbeat messages—directly or indirectly—to the rest of correct processes. Consider a run R of the algorithm in which \mathcal{C} processes are correct and let t be the time instant after which only $\mathcal{C} - 1$ unidirectional links carry messages forever. Consider now another run R' , identical to R until time t , and assume that a process q different from the leader, which is

correct in R , crashes at time t in R' . Observe that both executions R and R' are indistinguishable, and there is no way for the leader to know that q has crashed, and hence it will not stop sending messages to q . Since the number of correct processes in run R' is $\mathcal{C} - 1$, the algorithm should use only $\mathcal{C} - 2$ unidirectional links to carry messages forever, which contradicts the fact that the leader will not stop sending messages to q .

In [6], Aguilera et al. propose an algorithm implementing Ω such that eventually only f links carry messages forever, being f the maximum number of processes that can crash. They also show that in the crash-failure model no algorithm using fewer than f links exists. Hence, if $f = n - 1$ (as in our system model), Ω can be implemented with $n - 1$ links carrying messages forever, even if no process crashes, i.e., $\mathcal{C} = n$. However, the algorithm of [6] uses always $n - 1$ links, independently of the actual number of correct processes \mathcal{C} . As we will see, the algorithms proposed in this paper, besides implementing $\diamond\mathcal{P}$, dynamically adapt the number of links used forever to the actual number of correct processes.

Similar results can be deduced from the work by Fernández et al. in [9]. They study the minimal system conditions to implement unreliable failure detectors, and focus on the set *Reach* of correct processes that can reach all correct processes via exclusively eventually timely links and other correct processes. They show that $\diamond\mathcal{P}$ cannot be implemented if *Reach* does not contain all the correct processes. Similarly, they show that $\diamond\mathcal{S}$ (and hence Ω) cannot be implemented if *Reach* does not contain at least one correct process. In both cases, the subgraph formed by correct processes and eventually timely links in their system model must contain at least \mathcal{C} arcs (e.g., in a ring

topology), with $\mathcal{C} \leq n - 1$ if at least one process crashes. In terms of our system model, these arcs correspond to our \mathcal{C} links carrying messages forever.

4. Communication-Optimal Implementations of $\diamond\mathcal{P}$

In this section, we introduce two different approaches to the design of communication-optimal failure detection algorithms implementing $\diamond\mathcal{P}$. The approaches differ in how failure suspicions are managed. The first one uses an eager strategy in order to get low detection latencies. The second one is much more conservative in order to have a low sporadic communication overhead.

The first approach is based in every process consistently managing a local balance of suspicions and refutations for any other process. When a process p suspects another process q , p broadcasts a suspicion to every process, including q . If q has not failed, upon reception of that suspicion it will broadcast in the same way a refutation. Suspicions increment the corresponding balance, while refutations decrement it. With this strategy, eventually every correct process will permanently have a positive balance for every incorrect process, and a zero balance for every correct process. Observe that a reliable broadcast of suspicions and refutations is required in order to have consistent balances.

The alternative approach to the global spread of suspicions and refutations is to let a process to manage only suspicions in its *neighborhood* in the ring. In this approach, a process p will be in charge of the detection of incorrect processes between p 's correct predecessor in the ring and p . The ring arrangement is used to propagate information about failures, piggybacked on

periodic heartbeat messages, to all processes in a lazy way.

4.1. An Algorithm Using Reliable Broadcast

In this section, we propose a first communication-optimal implementation of $\diamond\mathcal{P}$ that uses Reliable Broadcast. In the algorithm, each process sends heartbeats to its successor in the ring, and monitors its predecessor by hearing heartbeats from it. Algorithm 2 presents the algorithm in detail, which uses a $Balance_p$ variable for every process p , accounting suspicions and refutations for every process. If $Balance_p(q) > 0$ with $q \neq p$, then p suspects q ; else, q is trusted by p . As we will see, $Balance_p$ provides the properties of $\diamond\mathcal{P}$. Every process p starts sending periodically an $(ALIVE, p)$ message to its successor in the ring, denoted by the variable $succ_p$ (Task 1). Also, every process p waits for periodical $(ALIVE, pred_p)$ messages from its predecessor in the ring, denoted by the variable $pred_p$. If p does not receive such a message on a specific time-out interval of $\Delta_p(pred_p)$, then p suspects that $pred_p$ has crashed, and R-broadcasts a $(SUSPICION, p, pred_p)$ message (Task 2), as shown in Figure 2a (p_1 suspects p_7). In Task 3, when p R-delivers a $(SUSPICION, q, r)$ message, p increments $Balance_p(r)$ and calls the $update_pred_and_succ()$ procedure. Besides this, if $r = p$, i.e., p has been erroneously suspected by q , p R-broadcasts a $(REFUTATION, p)$ message (Figure 2b). In Task 4, when p R-delivers a $(REFUTATION, q)$ message, p decrements $Balance_p(q)$, increments $\Delta_p(q)$, and calls the $update_pred_and_succ()$ procedure. Variables $pred_p$ and $succ_p$ are updated from $Balance_p$ to the nearest predecessor and the nearest successor in the ring having a non-positive balance respectively.⁵

⁵Here we informally use the terms *nearest predecessor* (or *nearest successor*) of a process p to denote the first process preceding (or succeeding) p following the ring arrangement

If all the components of the $Balance_p$ vector are positive, then p sets both $pred_p$ and $succ_p$ to p .




Figure 2: Sporadic communication in the communication-optimal Algorithm 2. Note that relayed messages caused by Reliable Broadcast are not shown.

Correctness Proof

We show now that Algorithm 2 implements a failure detector of class $\diamond\mathcal{P}$ and that it is communication optimal. In the proof, we consider that all the time instants are after all the incorrect processes have already crashed, and all the messages they have sent/R-broadcast before crashing have already been received/R-delivered. We start making the following observations.

Observation 1. $\forall p \in correct$, eventually and permanently $Balance_p(p) = 0$. This derives from the following: (1) initially $Balance_p(p) = 0$, (2) for every $(SUSPICION, -, p)$ message that p R-delivers in Task 3, eventually p R-delivers in Task 4 a $(REFUTATION, p)$ message that compensates the previous increment of $Balance_p(p)$, and (3) eventually every correct process

and fitting a particular condition.

Algorithm 2: Communication-optimal $\diamond\mathcal{P}$ using Reliable Broadcast.

```

{Every process  $p$  executes the following}

Procedure update_pred_and_succ()
  if  $\forall r : Balance_p(r) > 0$  then
     $pred_p \leftarrow p$ 
     $succ_p \leftarrow p$ 
  else
     $pred_p \leftarrow p$ 's nearest predecessor  $r$  in the ring such that  $Balance_p(r) \leq 0$ 
     $succ_p \leftarrow p$ 's nearest successor  $r$  in the ring such that  $Balance_p(r) \leq 0$ 

 $pred_p \leftarrow pred(p)$            { $p$ 's estimation of its nearest correct predecessor in the ring}
 $succ_p \leftarrow succ(p)$        { $p$ 's estimation of its nearest correct successor in the ring}
forall  $q \in \Pi$  do
   $\Delta_p(q) \leftarrow$  default time-out interval   { $\Delta_p(q)$  denotes the duration of  $p$ 's time-out for  $q$ }
   $Balance_p(q) \leftarrow 0$ 

cobegin
  || Task 1: repeat periodically
  || if  $succ_p \neq p$  then
  ||    $\lfloor$  send (ALIVE,  $p$ ) to  $succ_p$ 
  ||
  || Task 2: repeat periodically
  || if  $\left( \begin{array}{l} pred_p \neq p \text{ and } p \text{ did not receive } (ALIVE, pred_p) \\ \text{during the last } \Delta_p(pred_p) \text{ ticks of } p\text{'s clock} \end{array} \right)$  then   {time-out}
  ||    $\lfloor$  r-broadcast (SUSPICION,  $p$ ,  $pred_p$ )
  ||
  || Task 3: when r-deliver (SUSPICION,  $q$ ,  $r$ )
  ||    $Balance_p(r) \leftarrow Balance_p(r) + 1$ 
  ||   update_pred_and_succ()
  ||   if  $r = p$  then
  ||      $\lfloor$  r-broadcast (REFUTATION,  $p$ )
  ||
  || Task 4: when r-deliver (REFUTATION,  $q$ )
  ||    $Balance_p(q) \leftarrow Balance_p(q) - 1$ 
  ||    $\Delta_p(q) \leftarrow \Delta_p(q) + 1$            {not needed if  $q = p$ }
  ||   update_pred_and_succ()
coend

```

stabilizes with its correct predecessor in the ring, after which it stops R-broadcasting suspicions.

Observation 2. $\forall p$, if $\text{pred}_p = q$ with $q \neq p$ then $\text{Balance}_p(q) \leq 0$. Also, if $\text{succ}_p = r$ with $r \neq p$ then $\text{Balance}_p(r) \leq 0$. This derives directly from the fact that both pred_p and succ_p are only updated by p inside the procedure `update_pred_and_succ()`.

Observation 3. For every pair of correct processes p, q , when no more failure suspicions occur, $\text{Balance}_p(r) = \text{Balance}_q(r)$ for every process r . By the properties of Reliable Broadcast, both p and q R-deliver the same set of $(\text{SUSPICION}, -, r)$ and $(\text{REFUTATION}, r)$ messages. Consequently, both p and q apply the same modifications to $\text{Balance}_p(r)$ and $\text{Balance}_q(r)$ respectively.

Lemma 1. For every pair of consecutive correct processes q, p in the ring, eventually p stops R-broadcasting $(\text{SUSPICION}, p, q)$ messages.

PROOF. The proof is by contradiction. Assume that p R-broadcasts $(\text{SUSPICION}, p, q)$ messages infinitely often. Since both p and q are correct, for each message $(\text{SUSPICION}, p, q)$ q will R-broadcast a $(\text{REFUTATION}, q)$ message that will be R-delivered by p . Upon R-delivery, p will increment $\Delta_p(q)$. Since the communication link between q and p is eventually timely, eventually $\Delta_p(q)$ will reach the unknown bound on message transmission times, after which p will receive an (ALIVE, q) message always before $\Delta_p(q)$ expires, and p will no more suspect q in Task 2. This contradicts the fact that p suspects q infinitely often.

Lemma 2. For every pair of non-consecutive correct processes q, p in the ring, eventually p stops R-broadcasting $(\text{SUSPICION}, p, q)$ messages.

PROOF. By Lemma 1, eventually p will permanently monitor another correct process r , being r its correct predecessor in the ring. After that, p will never R-broadcast any $(SUSPICION, p, q)$ message any more.

Lemma 3. *For every pair of correct processes p, q , eventually and permanently $Balance_p(q) = 0$.*

PROOF. Follows from Lemma 1 and Lemma 2, and the fact that, being initially $Balance_p(q) = 0$, by the algorithm p R-delivers the same number of $(SUSPICION, -, q)$ and $(REFUTATION, q)$ messages, compensating the increment and decrement operations over $Balance_p(q)$.

Lemma 4. *For every incorrect process q , eventually and permanently $Balance_p(q) > 0$ for every correct process p .*

PROOF. Note that after q crashes it will not be able to R-broadcast any message $(REFUTATION, q)$. Also, at least process r , being r the correct successor of q in the ring, will eventually R-broadcast a $(SUSPICION, r, q)$ message that q will not refute, and consequently $Balance_r(q) > 0$ permanently. Then, by Observation 3, $Balance_p(q) > 0$ for every correct process p .

Lemma 5. *Eventually, for every correct process p , $pred_p$ will be permanently set to p 's correct predecessor in the ring, and $succ_p$ will be permanently set to p 's correct successor in the ring.*

PROOF. There are two cases to consider. If $\mathcal{C} = 1$, then by definition both p 's correct predecessor in the ring and p 's correct successor in the ring are p

itself, and the lemma follows directly from Lemma 4, Observation 1 and the procedure *update_pred_and_succ()*. Otherwise, i.e., if $\mathcal{C} > 1$, the lemma follows directly from Lemma 3, Lemma 4, Observation 2 and the procedure *update_pred_and_succ()*.

Theorem 3. *Algorithm 2 implements a failure detector of class $\diamond\mathcal{P}$.*

PROOF. From Lemma 3 and Lemma 4, for every correct process p , eventually and permanently $Balance_p(q) = 0$ for every $q \in correct$, and $Balance_p(r) > 0$ for every $r \in crashed$. The rule “if $Balance_p(q) > 0$, then p suspects q ; else, p does not suspect q ” provides the properties of strong completeness and eventual strong accuracy of $\diamond\mathcal{P}$.

Theorem 4. *Algorithm 2 is communication-optimal, i.e., eventually only \mathcal{C} links carry messages forever.*

PROOF. From Lemma 5, for every correct process p , eventually and permanently $succ_p$ will be set to p 's correct successor in the ring and, by Task 1, p will send (*ALIVE*, p) messages to it forever. No other periodical messages will be sent. Furthermore, since no more suspicions will occur, no new *SUSPICION* (and hence *REFUTATION*) messages will be broadcast. Thus, if there are \mathcal{C} correct processes in the system, just a number of \mathcal{C} unidirectional links will be permanently used.

Observe that if there is just one correct process in the system, i.e., $\mathcal{C} = 1$, Algorithm 2 eventually uses no links, by an optimization introduced in Task 1. Hence, when $\mathcal{C} = 1$ both $\diamond\mathcal{P}$ and Ω can be implemented using 0 links carrying messages forever.

Finally, although we have initially assumed in the system model that all the communication links were reliable and eventually timely, in a given execution of Algorithm 2 it is sufficient that this behaviour applies only to the \mathcal{C} links that eventually form the ring of correct processes, i.e., the links from every correct process to its correct successor in the ring. The rest of links can be asynchronous and/or lossy.

4.2. An Algorithm Using One-to-One Local Communication

In this section, we present a communication-optimal $\diamond\mathcal{P}$ algorithm that uses only one-to-one local communication to manage suspicions. In the algorithm, a process p will be in charge of the detection of incorrect processes between p 's correct predecessor in the ring and p . Assuming that simple heartbeat messages circulate around the ring, this strategy only gives *weak* completeness and eventual strong accuracy, and hence the resulting failure detector will be of the class $\diamond\mathcal{Q}$. Henceforth, a further transformation is needed to get a failure detector of the class $\diamond\mathcal{P}$. In this way, the approach we follow to design the algorithm is incremental: first we present the algorithm implementing $\diamond\mathcal{Q}$ and prove its correctness, and next we give a simple transformation into $\diamond\mathcal{P}$ which preserves communication optimality and low sporadic communication overhead.

4.2.1. Implementing $\diamond\mathcal{Q}$

Algorithm 3 presents a communication-optimal implementation of $\diamond\mathcal{Q}$. Every process p has a local set of suspected processes, L_p , and two variables, $pred_p$ and $succ_p$, denoting respectively the process that p is monitoring and the process to which p is periodically sending heartbeat messages (*ALIVE*, p)

Algorithm 3: Communication-optimal $\diamond Q$ using local one-to-one communication.

{Every process p executes the following}

Procedure *udpate_pred_and_succ()*

```

   $pred_p \leftarrow p$ 's nearest predecessor  $r$  in the ring such that  $r \notin L_p$ 
   $succ_p \leftarrow p$ 's nearest successor  $r$  in the ring such that  $r \notin L_p$ 
  if  $pred_p \neq p$  then
     $L_p \leftarrow \{pred_p, \dots, succ_p\} - \{pred_p, p, succ_p\}$ 

   $pred_p \leftarrow pred(p)$  { $p$ 's estimation of its nearest correct predecessor in the ring}
   $succ_p \leftarrow succ(p)$  { $p$ 's estimation of its nearest correct successor in the ring}
  forall  $q \in \Pi$  do
     $\Delta_p(q) \leftarrow$  default time-out interval { $\Delta_p(q)$  denotes the duration of  $p$ 's time-out for  $q$ }
     $L_p \leftarrow \emptyset$  { $L_p$  provides the properties of  $\diamond Q$ }

```

cobegin

```

  || Task 1: repeat periodically
  || if  $succ_p \neq p$  then
  ||    $\lfloor$  send (ALIVE,  $p$ ) to  $succ_p$ 

  || Task 2: repeat periodically
  || if  $\left( \begin{array}{l} pred_p \neq p \text{ and } p \text{ did not receive } (ALIVE, pred_p) \\ \text{during the last } \Delta_p(pred_p) \text{ ticks of } p\text{'s clock} \end{array} \right)$  then {time-out}
  ||    $L_p \leftarrow L_p \cup \{pred_p\}$  { $p$  suspects  $pred_p$  has crashed}
  ||   send (SUSPICION,  $p$ ) to  $pred_p$ 
  ||   update_pred_and_succ()

  || Task 3: when receive (SUSPICION,  $q$ ) from some  $q$ 
  ||    $L_p \leftarrow L_p \cup \{p, \dots, q\} - \{p, q\}$ 
  ||   update_pred_and_succ()
  ||   forall  $r \in \{p, \dots, q\} - \{p, q\}$  do
  ||      $\lfloor$  send (PROBE,  $p$ ) to  $r$ 
  ||   send (ALIVE,  $p$ ) to  $q$ 

  || Task 4: when receive (ALIVE,  $q$ ) from some  $q$ 
  || if  $q \in L_p$  then
  ||    $L_p \leftarrow L_p - \{q\}$ 
  ||   update_pred_and_succ()
  ||    $\Delta_p(q) \leftarrow \Delta_p(q) + 1$ 

  || Task 5: when receive (PROBE,  $q$ ) from some  $q$ 
  ||   send (ALIVE,  $p$ ) to  $q$ 

```

coend

by Task 1. When a process p suspects by Task 2 the process it is monitoring, $pred_p$, p includes $pred_p$ in L_p , sends a suspicion message ($SUSPICION, p$) to $pred_p$ (in Figure 3a, p_1 suspects p_6), and updates $pred_p$ (and $succ_p$ if required) accordingly. Observe that Task 2 does not include any explicit mechanism for p to tell its new predecessor to start sending heartbeats to it. The new predecessor of p will have to be suspected once by p in order to set its successor to p by Task 3, as we will explain next.

If a suspected process p is correct or has not crashed yet, when it receives ($SUSPICION, q$) by Task 3, p will suspect every process from p to q (both excluded), since all of them have been also suspected by q . Consequently, process q becomes the new successor of p , and hence, if p does not crash, q will receive periodical ($ALIVE, p$) messages from p , as shown in Figure 3b. Furthermore, a sporadic ($PROBE, p$) message is sent by p to every process r from p to q (both excluded), in order to know if r has actually crashed or not (also shown in Figure 3b, in which p_6 probes p_7 and p_8). When a process p receives a ($PROBE, q$) message, it just sends an ($ALIVE, p$) message to q , in order to give q the opportunity to set $succ_q$ (and exceptionally $pred_q$) to p (in Figure 3c, p_7 and p_8 notify p_6 that they are alive).

On the reception of an ($ALIVE, q$) message coming from a process $q \in L_p$, by Task 4 a process p will remove q from L_p , update $pred_p$ (and $succ_p$ if required), and increment the time-out interval with respect to q , $\Delta_p(q)$.

The $PROBE$ messages used in this algorithm avoid a process to send periodically $ALIVE$ messages to crashed processes, a scenario that can happen in communication-efficient algorithms (see Figure 1, where processes p_7 and p_8 have crashed). Hence, the proposed probing mechanism is key to get

communication optimality.

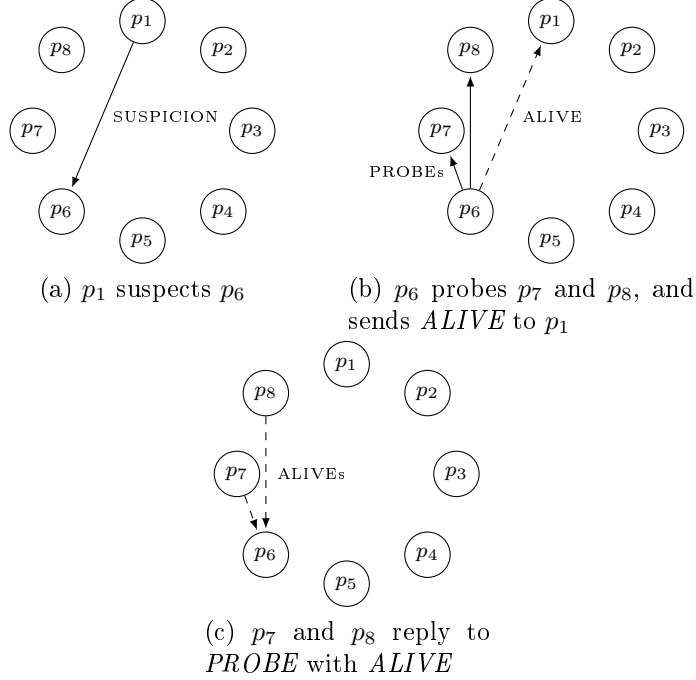


Figure 3: Sporadic communication in the communication-optimal Algorithm 3.

Correctness Proof

We show now that Algorithm 3 implements a failure detector of class $\diamond\mathcal{Q}$ and that it is communication-optimal. Given any process p , we denote by $corr_pred_p$ the correct predecessor of p in the ring. Similarly, we denote by $corr_succ_p$ the correct successor of p in the ring. The key of the proof is to show that eventually and permanently $pred_p = corr_pred_p$ and $succ_p = corr_succ_p$ for every correct process p . In other words, the ring stabilizes in terms of both the $pred$ and $succ$ variables of processes, which guarantees the correct construction of the sets L_p of suspected processes.

We assume that every task is executed as a critical section. We start by making the following observations:

Observation 4. $p \notin L_p$ permanently for every process p .

Observation 5. $L_p = \{pred_p, \dots, succ_p\} - \{pred_p, p, succ_p\}$ permanently after the execution of any task for every process p .

Observe that whenever L_p is modified by Task 2, Task 3 or Task 4 of p , $pred_p$ and $succ_p$ (and L_p itself) are updated accordingly by the procedure `update_pred_and_succ()`.

Observation 6. Whenever a process q is included by a correct process p in L_p , p will send a message (of type *SUSPICION* in Task 2 or type *PROBE* in Task 3) to q , and, if q is correct, p will eventually receive an $(ALIVE, q)$ message sent by Task 3 or Task 5 of q .

For the rest of the proof we will assume that any time instant t considered is larger than a time t_{base} that occurs after the stabilization time GST (i.e., $t_{base} > GST$), after every incorrect process has crashed, and after all messages sent by incorrect processes have been received. Note that this eventually happens. Hence, any new message received has necessarily been sent by a correct process.

Lemma 6. For every correct process p , eventually and permanently $pred_p = corr_pred_p$.

PROOF. For a correct process p , let $pred_p = r$ such that $r \neq corr_pred_p$. For a given time in the execution of the algorithm, one of the following cases applies:

Case 1: assume first that $r \in \{corr_pred_p, \dots, p\} - \{corr_pred_p, p\}$. Observe that, with this assumption, r is by definition not correct, as well as any other process in that range. Therefore, r will be eventually included in L_p (by Task 2 or Task 3 of p), and, since r has already crashed, it will not be able to send an $(ALIVE, r)$ message to p , remaining in L_p forever.

Case 2: assume now that $r \notin \{corr_pred_p, \dots, pred(p)\}$. By Observation 5, $corr_pred_p \in L_p$. Observe that $corr_pred_p$ could have been included in L_p by Task 2 or by Task 3 of p . Since $corr_pred_p$ is by definition correct, by Observation 6 eventually p will receive and $(ALIVE, corr_pred_p)$ message. At this point, $pred_p$ will be set to $corr_pred_p$.

Since every time $corr_pred_p$ has been suspected by p , $\Delta_p(corr_pred_p)$ is incremented by Task 4 of p , and since the communication link between $corr_pred_p$ and p is eventually timely, eventually $\Delta_p(corr_pred_p)$ will reach the unknown bound on message transmission times, after which p will receive the periodical $(ALIVE, corr_pred_p)$ messages always before $\Delta_p(corr_pred_p)$ expires and $corr_pred_p$ will not be suspected by p anymore.

Lemma 7. *For every correct process p , eventually and permanently $succ_p = corr_succ_p$.*

PROOF. The proof is by contradiction. By Lemma 6, eventually and permanently $pred_{corr_succ_p} = p$. Assume that $succ_p \neq corr_succ_p$. Since p is not sending by Task 1 periodical $(ALIVE, p)$ messages to $corr_succ_p$, then by Task 2 $corr_succ_p$ will eventually suspect p and modify $pred_{corr_succ_p}$, which contradicts Lemma 6.

Theorem 5. *Algorithm 3 implements a failure detector of class $\diamond Q$.*

PROOF. From Lemmas 6 and 7, for every correct process p , eventually p will be in the stable ring formed by correct processes. Otherwise, p is incorrect, and by Lemma 6 eventually and permanently $pred_{corr_succ_p} = corr_pred_p$. By Observation 5, p will eventually and permanently be included in $L_{corr_succ_p}$. As a consequence, eventually and permanently p will be included in the set of suspected processes of some correct process. This provides the weak completeness property of $\diamond Q$.

By Observation 4, p is never included in L_p . Once the ring has stabilized, no correct process is included (in Task 2) in the set L_p of any correct process p . Hence, once the ring has stabilized, no correct process will be present in any set of suspected processes. This provides the eventual strong accuracy property of $\diamond Q$.

Theorem 6. *Algorithm 3 is communication-optimal, i.e., eventually only \mathcal{C} links carry messages forever.*

PROOF. From Lemma 7, for every correct process p , eventually and permanently $succ_p$ will be set to p 's correct successor in the ring and, by Task 1, p will send (*ALIVE*, p) messages to it forever. No other periodical messages will be sent. Furthermore, since no more suspicions will occur, no new *SUSPICION* (and hence sporadic *PROBE* or *ALIVE*) messages will be sent. Thus, if there are \mathcal{C} correct processes in the system, just a number of \mathcal{C} unidirectional links will be permanently used.

A similar reasoning as the one made for the previous algorithm can be made for Algorithm 3 regarding the minimal assumptions on communication reliability and synchrony required by the algorithm. In this case, in a given

execution of Algorithm 3 it is sufficient that the links from every correct process to both its correct successor and predecessor in the ring are reliable and eventually timely, i.e., $2\mathcal{C}$ links. The rest of links can be asynchronous and/or lossy.

4.2.2. Transforming $\diamond\mathcal{Q}$ into $\diamond\mathcal{P}$

In this section we present a transformation of the $\diamond\mathcal{Q}$ algorithm of the previous section into $\diamond\mathcal{P}$. The transformation preserves the communication optimality and low sporadic communication overhead of Algorithm 3.

Algorithm 4, which implements $\diamond\mathcal{P}$, is obtained from Algorithm 3 by adding a set G_p to every process p . G_p is included into the *ALIVE* messages sent by p . In the algorithm, additions and removals of processes to L_p are also applied to G_p . To build G_p from the set G_{pred_p} received in Task 4, process p adds the processes between $pred_p$ and p to G_{pred_p} (both $pred_p$ and p excluded). In other words, p relies on its predecessor for suspicions beyond its directly monitored neighborhood.

Correctness Proof

We show now that Algorithm 4 implements a failure detector of class $\diamond\mathcal{P}$ and that it is also communication-optimal. The key of the proof is to show that Algorithm 4 is a transformation of Algorithm 3 into $\diamond\mathcal{P}$ preserving communication optimality. Observe that in Task 2 and Task 3 exactly the same operations are applied to L_p and to G_p . Finally, in Task 4, the global set of suspected processes G_p is built by adding to G_{pred_p} the processes between $pred_p$ and p (both $pred_p$ and p excluded).

Observation 7. $p \notin G_p$ permanently for every process p .

Algorithm 4: Communication-optimal $\diamond\mathcal{P}$ using local one-to-one communication.

```

{Every process  $p$  executes the following}
Procedure update_pred_and_succ()
   $pred_p \leftarrow p$ 's nearest predecessor  $r$  in the ring such that  $r \notin L_p$ 
   $succ_p \leftarrow p$ 's nearest successor  $r$  in the ring such that  $r \notin L_p$ 
  if  $pred_p \neq p$  then
     $L_p \leftarrow \{pred_p, \dots, succ_p\} - \{pred_p, p, succ_p\}$ 

 $pred_p \leftarrow pred(p)$                                 { $p$ 's estimation of its nearest correct predecessor in the ring}
 $succ_p \leftarrow succ(p)$                             { $p$ 's estimation of its nearest correct successor in the ring}
forall  $q \in \Pi$  do
   $\Delta_p(q) \leftarrow$  default time-out interval        { $\Delta_p(q)$  denotes the duration of  $p$ 's time-out for  $q$ }
   $L_p \leftarrow \emptyset$                                { $L_p$  provides the properties of  $\diamond\mathcal{Q}$ }
   $G_p \leftarrow \emptyset$                                { $G_p$  provides the properties of  $\diamond\mathcal{P}$ }

cobegin
  || Task 1: repeat periodically
  || if  $succ_p \neq p$  then
  ||    $\lfloor$  send (ALIVE,  $p$ ,  $G_p$ ) to  $succ_p$ 
  ||
  || Task 2: repeat periodically
  || if (  $pred_p \neq p$  and  $p$  did not receive (ALIVE,  $pred_p$ ,  $-$ ) ) then                                {time-out}
  ||   ( during the last  $\Delta_p(pred_p)$  ticks of  $p$ 's clock )
  ||    $L_p \leftarrow L_p \cup \{pred_p\}$                                                                 { $p$  suspects  $pred_p$  has crashed}
  ||    $G_p \leftarrow G_p \cup \{pred_p\}$ 
  ||   send (SUSPICION,  $p$ ) to  $pred_p$ 
  ||    $\lfloor$  update_pred_and_succ()
  ||
  || Task 3: when receive (SUSPICION,  $q$ ) from some  $q$ 
  ||    $L_p \leftarrow L_p \cup \{p, \dots, q\} - \{p, q\}$ 
  ||    $G_p \leftarrow G_p \cup \{p, \dots, q\} - \{p, q\}$ 
  ||   update_pred_and_succ()
  ||   forall  $r \in \{p, \dots, q\} - \{p, q\}$  do
  ||      $\lfloor$  send (PROBE,  $p$ ) to  $r$ 
  ||   send (ALIVE,  $p$ ,  $G_p$ ) to  $q$ 
  ||
  || Task 4: when receive (ALIVE,  $q$ ,  $G_q$ ) from some  $q$ 
  ||   if  $q \in L_p$  then
  ||      $L_p \leftarrow L_p - \{q\}$ 
  ||     update_pred_and_succ()
  ||      $\Delta_p(q) \leftarrow \Delta_p(q) + 1$ 
  ||   if  $q = pred_p$  then
  ||      $\lfloor$   $G_p \leftarrow (G_q \cup \{pred_p, \dots, p\}) - \{pred_p, p\}$ 
  ||
  || Task 5: when receive (PROBE,  $q$ ) from some  $q$ 
  ||    $\lfloor$  send (ALIVE,  $p$ ,  $G_p$ ) to  $q$ 

coend

```

Theorem 7. *Algorithm 4 implements a failure detector of class $\diamond\mathcal{P}$.*

PROOF. Observe that, by the construction of the set G_p in Tasks 2, 3, and 4,

G_p is equivalent to L_p in the range $\{pred_p, \dots, p\}$. By Lemmas 6 and 7, and by Task 4 of p , eventually G_p will include every incorrect process, providing the strong completeness property of $\diamond\mathcal{P}$. By Observation 7 and again by Lemmas 6 and 7 and Task 4 of p , no correct process will be eventually included in G_p , preserving the eventual strong accuracy property of $\diamond\mathcal{P}$.

Theorem 8. *Algorithm 4 is communication-optimal, i.e., eventually only \mathcal{C} links carry messages forever.*

PROOF. Follows directly from the fact that no additional message is sent in Algorithm 4 with respect to Algorithm 3.

5. Performance Analysis

In Subsection 1.2 we have defined a set of properties to be taken into account when designing failure detectors, namely communication efficiency, low sporadic overhead and communication locality. In the light of these properties, in this section we analyze the performance of the communication-optimal $\diamond\mathcal{P}$ algorithms. We also include in the analysis Chandra-Toueg’s algorithm [1] as a reference, as well as the simplest of the communication-efficient algorithms in [21]. The goal is to highlight the strengths and weakness of each approach and set the trade-offs to be considered in the implementation of failure detectors for particular scenarios. Specifically, we focus on the following performance parameters:

- Communication efficiency/optimality. Number of unidirectional links that carry messages forever.

- Sporadic communication overhead. Number of extra messages exchanged to manage an erroneous suspicion.
- Detection latency (or system responsiveness). Upon a crash, time elapsed until every alive process permanently suspects the crashed process.
- Communication locality. Scope of the messages. An algorithm exhibits (periodic or sporadic) communication locality when processes send (periodic or sporadic) messages mainly to their neighborhood.

Table 2 summarizes the performance of the algorithms. Note that sporadic communication locality is not applicable to the algorithms in [1, 21], since they do not have sporadic communication at all (sporadic overhead is 0). Observe also that Algorithm 4, besides being communication-optimal, has a low sporadic overhead due to its one-to-one communication pattern, and has both periodic and sporadic communication locality. This is at the cost of a linear detection latency, in contrast to the uniform detection latency of the algorithms using a one-to-all communication pattern.

Algorithm	Comm. efficiency (# links used forever)	Sporadic overhead (# extra messages)	Detection latency	Comm. locality	
				Periodic	Sporadic
Chandra-Toueg [1]	$\mathcal{C}(n-1)$	0	T_h	No	—
Comm.-efficient [21]	n	0	$\mathcal{C} T_h$	Yes	—
Comm.-optimal Alg. 2	\mathcal{C}	$2n^2$	T_h	Yes	No
Comm.-optimal Alg. 4	\mathcal{C}	$2l$	$\mathcal{C} T_h$	Yes	Yes

Table 2: Performance analysis of $\diamond\mathcal{P}$ algorithms. Sporadic overhead and detection latency figures are approximate (l denotes the cardinality of the local set of suspected processes, usually $l \ll n$).

In Table 2, detection latency is estimated on the basis of the term T_h , denoting the mean delay of sending a new, periodic, heartbeat message, by which the information about suspected processes is relayed to the successor process in the ring in the algorithm in [21] and in Algorithm 4. Assuming that the transmission delay of any message is negligible with respect to T_h , the value of T_h can be estimated as one half of the periodicity of the task that sends periodic heartbeat messages (Task 1 in the algorithms proposed in this paper).

The linear detection latency in Algorithm 4 is due to the fact that the information about a new suspicion has to circulate around the whole ring to reach every alive process. Nevertheless, a mechanism to speed-up the detection latency, at the price of losing sporadic communication locality, can be easily added. The mechanism consists in the suspecting process p directly sending the suspicion message not only to the suspected process but also to a subset of processes $\Lambda \subset \Pi$, i.e., adding shortcuts in the ring. In general, if $k = |\Lambda|$, while the communication overhead is incremented in at most $2k$ messages, assuming that shortcuts have been uniformly distributed, an estimation of the maximum detection latency results in $(\frac{n}{k+1})T_h$. In the limit, $\Lambda = \Pi - \{p\}$, and the detection latency is reduced to the minimum. The mechanism can be also applied to the communication-efficient algorithm in [21].

Figure 4 shows an example of the shortcut mechanism to improve responsiveness in Algorithm 4. Without shortcuts, a failure suspicion propagates by heartbeats and reaches all alive processes in approximately $\mathcal{C} T_h$ time. One or more shortcuts allow the suspicion to be propagated in parallel, reducing

the detection latency. Of course, in case the suspicion is erroneous, such a mechanism makes the error to be propagated faster in the system. A way of avoiding this apparent drawback consists in adding a complementary shortcut mechanism for the refutation of a recent, erroneous, suspicion. Observe that neither of these suspicion/refutation shortcut mechanisms compromises the optimality of Algorithm 4, since eventually, i.e., when the ring formed by correct processes stabilizes, they are not used any more.

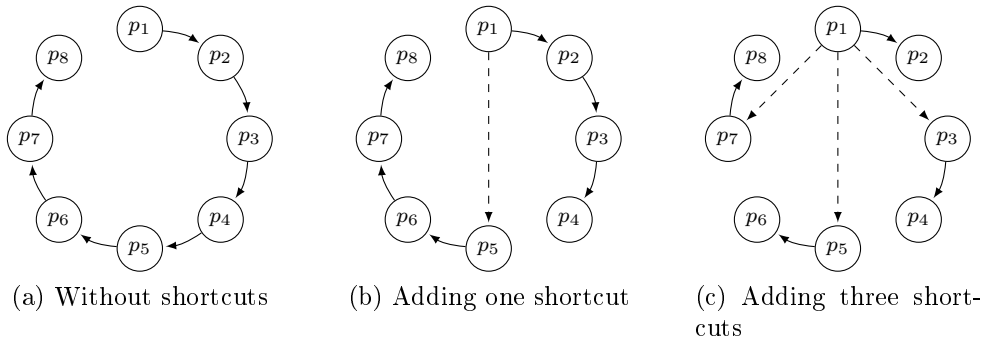


Figure 4: Example of the shortcut mechanism to improve responsiveness.

In general, a trade-off results between detection latency and communication overhead, being the latter either periodic or sporadic. It is worth noting that parameters related to communication overhead have been expressed in terms of the number of links and messages. Communication costs, however, strongly depend on network parameters as topology and link delays, specifically in wide area networks, resulting in not uniform message transmission delays. This fact can be exploited when using a ring arrangement as a pattern for communication. More precisely, communication costs associated to pairs of processes provide a hint for the initial ring configuration. Periodic heartbeat communication will take advantage of this. Furthermore, sporadic

communication overhead, which is inherent to communication-optimal algorithms, results in practice less harmful for algorithms with this kind of local communication than for algorithms with global communication. Finally, as we have shown, the ability of using shortcuts in the formers provide an interesting basis to fit trade-offs between sporadic communication overhead and detection latency.

6. Conclusions

In this paper, we have explored communication efficiency, a performance measure that refers to the number of unidirectional links that carry messages forever in an algorithm. We have shown that failure detector class $\diamond\mathcal{P}$ requires at least \mathcal{C} unidirectional links to carry messages forever, being \mathcal{C} the number of correct processes. Moreover, when at least one process crashes, \mathcal{C} links are also required for $\diamond\mathcal{S}$ and Ω . We have proposed two ring-based communication-optimal $\diamond\mathcal{P}$ algorithms. Since these algorithms use exactly \mathcal{C} unidirectional links to carry messages forever, it can be derived that communication optimality for $\diamond\mathcal{P}$ is achieved. Since $\diamond\mathcal{P}$ trivially implements Ω , communication optimality can be considered achieved also for $\diamond\mathcal{S}$ and Ω failure detectors.

One of the algorithms uses a Reliable Broadcast primitive to communicate suspicions and refutations, involving a quadratic number of messages. Since this can be a drawback in some scenarios, e.g., very large networks, we have proposed a second algorithm that uses exclusively one-to-one communication. This algorithm has some interesting properties which make it suitable for distributed applications deployed over wide area networks: (1) communication

optimality, i.e., only \mathcal{C} links carry messages forever, (2) low sporadic communication overhead to manage failure suspicions, i.e., the number of messages exchanged as a consequence of a suspicion is linear, and (3) communication locality, i.e., managing a failure suspicion at a process p only implies communicating with processes at p 's neighborhood in the ring. This is of particular interest in networks where communication costs between pairs of processes are not uniform. If the logical ring of processes is arranged regarding communication cost criteria, both communication overhead and delays of periodic heartbeats will be minimized.

The second algorithm admits a flexible implementation. In its basic form, i.e., using exclusively local communication, it exhibits a linear detection latency. However, a mechanism based on shortcuts can reduce it. Moreover, this mechanism can be enabled or disabled without affecting the correctness of the algorithm. As a consequence, the mechanism can be applied partially and be precisely tuned in order to fit trade-offs between network overhead and QoS parameters.

References

- [1] T. Chandra, S. Toueg, Unreliable failure detectors for reliable distributed systems, *Journal of the ACM* 43 (2) (1996) 225–267.
- [2] M. Pease, R. Shostak, L. Lamport, Reaching agreement in the presence of faults, *Journal of the ACM* 27 (2) (1980) 228–234.
- [3] T. Chandra, V. Hadzilacos, S. Toueg, The weakest failure detector for solving consensus, *Journal of the ACM* 43 (4) (1996) 685–722.

- [4] M. Aguilera, C. Delporte-Gallet, H. Fauconnier, S. Toueg, Stable leader election, in: Proceedings of the 15th International Symposium on Distributed Computing (DISC'2001), LNCS 2180, Springer-Verlag, Lisbon, Portugal, 2001, pp. 108–122.
- [5] M. Aguilera, C. Delporte-Gallet, H. Fauconnier, S. Toueg, On implementing Ω with weak reliability and synchrony assumptions, in: Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC'2003), Boston, Massachusetts, 2003, pp. 306–314.
- [6] M. Aguilera, C. Delporte-Gallet, H. Fauconnier, S. Toueg, Communication-efficient leader election and consensus with limited link synchrony, in: Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing (PODC'2004), St. John's, Newfoundland, Canada, 2004, pp. 328–337.
- [7] M. Aguilera, C. Delporte-Gallet, H. Fauconnier, S. Toueg, On implementing omega in systems with weak reliability and synchrony assumptions, Distributed Computing 21 (4) (2008) 285–314.
- [8] F. Chu, Reducing Ω to $\diamond\mathcal{W}$, Information Processing Letters 67 (6) (1998) 289–293.
- [9] A. Fernández, E. Jiménez, S. Arévalo, Minimal system conditions to implement unreliable failure detectors, in: Proceedings of the 12th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC'2006), University of California, Riverside, USA, 2006, pp. 63–72.

- [10] A. Fernández, E. Jiménez, M. Raynal, Eventual leader election with weak assumptions on initial knowledge, communication reliability, and synchrony, in: Proceedings of the IEEE International Conference on Dependable Systems and Networks (DSN'2006), Philadelphia, Pennsylvania, 2006, pp. 166–178.
- [11] E. Jiménez, S. Arévalo, A. Fernández, Implementing Unreliable Failure Detectors with Unknown Membership, *Information Processing Letters* 100 (2) (2006) 60–63.
- [12] M. Larrea, A. Fernández, S. Arévalo, Optimal implementation of the weakest failure detector for solving consensus, in: Proceedings of the 19th IEEE Symposium on Reliable Distributed Systems (SRDS'2000), Nuremberg, Germany, 2000, pp. 52–59.
- [13] A. Mostéfaoui, E. Mourgaya, M. Raynal, Asynchronous implementation of failure detectors, in: Proceedings of the IEEE International Conference on Dependable Systems and Networks (DSN'2003), San Francisco, California, 2003, pp. 351–360.
- [14] A. Mostéfaoui, E. Mourgaya, M. Raynal, C. Travers, A time-free assumption to implement eventual leadership, *Parallel Processing Letters* 16 (2) (2006) 189–208.
- [15] A. Mostéfaoui, D. Powell, M. Raynal, A hybrid approach for building eventually accurate failure detectors, in: Proceedings of the 10th IEEE Pacific Rim International Symposium on Dependable Computing (PRDC'2004), Papeete, Tahiti, 2004, pp. 57–65.

- [16] M. Larrea, A. Fernández, S. Arévalo, Eventually consistent failure detectors, *Journal of Parallel and Distributed Computing* 65 (3) (2005) 361–373.
- [17] M. Larrea, S. Arévalo, A. Fernández, Efficient algorithms to implement unreliable failure detectors in partially synchronous systems, in: *Proceedings of the 13th International Symposium on Distributed Computing (DISC'99)*, LNCS 1693, Springer-Verlag, Bratislava, 1999, pp. 34–48.
- [18] M. Larrea, A. Fernández, S. Arévalo, On the implementation of unreliable failure detectors in partially synchronous systems, *IEEE Transactions on Computers* 53 (7) (2004) 815–828.
- [19] M. Fischer, N. Lynch, M. Paterson, Impossibility of distributed consensus with one faulty process, *Journal of the ACM* 32 (2) (1985) 374–382.
- [20] C. Dwork, N. Lynch, L. Stockmeyer, Consensus in the presence of partial synchrony, *Journal of the ACM* 35 (2) (1988) 288–323.
- [21] M. Larrea, A. Lafuente, I. Soraluze, R. Cortiñas, J. Wieland, Designing efficient algorithms for the eventually perfect failure detector class, *Journal of Software* 2 (4) (2007) 1–11.
- [22] M. Hutle, Failure Detection in Sparse Networks, Ph.D. thesis, Vienna University of Technology, Wien (August 2005).
- [23] W. Chen, S. Toueg, M. Aguilera, On the quality of service of failure detectors, *IEEE Transactions on Computers* 51 (5) (2002) 561–580.

- [24] M. Hurfin, M. Raynal, A simple and fast asynchronous consensus protocol based on a weak failure detector, *Distributed Computing* 12 (4) (1999) 209–223.
- [25] A. Mostéfaoui, M. Raynal, Solving consensus using Chandra-Toueg’s unreliable failure detectors: a general quorum-based approach, in: *Proceedings of the 13th International Symposium on Distributed Computing (DISC’99)*, LNCS 1693, Springer-Verlag, Bratislava, 1999, pp. 49–63.
- [26] A. Mostéfaoui, M. Raynal, Leader-based consensus, *Parallel Processing Letters* 11 (1) (2001) 95–107.
- [27] A. Schiper, Early consensus in an asynchronous system with a weak failure detector, *Distributed Computing* 10 (3) (1997) 149–157.
- [28] F. Greve, M. Hurfin, R. Macedo, M. Raynal, Consensus based on strong failure detectors: A time and message-efficient protocol, in: *Proceedings of the 15th Parallel and Distributed Processing Workshop (IPDPS’2000)*, LNCS 1800, Springer-Verlag, Cancun, Mexico, 2000, pp. 1258–1265.
- [29] R. Guerraoui, M. Raynal, The information structure of indulgent consensus, *IEEE Transactions on Computers* 53 (4) (2004) 453–466.
- [30] L. Lamport, The part-time parliament, *ACM Transactions on Computer Systems* 16 (2) (1998) 133–169.
- [31] R. Guerraoui, M. Kapalka, P. Kouznetsov, The weakest failure detector to boost obstruction-freedom, in: *Proceedings of the 20th Interna-*

tional Symposium on Distributed Computing (DISC'2006), LNCS 4167, Springer-Verlag, Stockholm, Sweden, 2006, pp. 399–412.

- [32] W. Wu, J. Cao, J. Yang, M. Raynal, A hierarchical consensus protocol for mobile ad hoc networks, in: Proceedings of the 14th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP'2006), IEEE Computer Society, Montbeliard-Sochaux, France, 2006, pp. 64–72.
- [33] V. Hadzilacos, S. Toueg, Fault-tolerant broadcasts and related problems, in: S. J. Mullender (Ed.), Distributed Systems, 2nd Edition, Addison-Wesley, 1993, Ch. 5, pp. 97–146.