

Failure Detectors in Homonymous Distributed Systems (with an Application to Consensus)

Sergio Arévalo^a, Antonio Fernández Anta^b, Damien Imbs^c, Ernesto Jiménez^a, Michel Raynal^d

^a*EUI, Universidad Politécnica de Madrid, 28031 Madrid, Spain*

^b*Institute IMDEA Networks, 28918 Madrid, Spain*

^c*Instituto de Matemáticas, UNAM, D.F. 04510, México*

^d*Institut Universitaire de France and IRISA, Université de Rennes I, Campus de Beaulieu, 35042 Rennes, France*

Abstract

This paper is on homonymous distributed systems where processes are prone to crash failures and have no initial knowledge of the system membership (“homonymous” means that several processes may have the same identifier). New classes of failure detectors suited to these systems are first defined. Among them, the classes $H\Omega$ and $H\Sigma$ are introduced that are the homonymous counterparts of the classes Ω and Σ , respectively. (Recall that the pair $\langle\Omega, \Sigma\rangle$ defines the weakest failure detector to solve consensus.) Then, the paper shows how $H\Omega$ and $H\Sigma$ can be implemented in homonymous systems without membership knowledge (under different synchrony requirements). Finally, two algorithms are presented that use these failure detectors to solve consensus in homonymous asynchronous systems where there is no initial knowledge of the membership. One algorithm solves consensus with $\langle H\Omega, H\Sigma\rangle$, while the other uses only $H\Omega$, but needs a majority of correct processes.

Observe that the systems with unique identifiers and anonymous systems are extreme cases of homonymous systems from which follows that all these results also apply to these systems. Interestingly, the new failure detector class $H\Omega$ can be implemented with partial synchrony¹, while the analogous class $A\Omega$ defined for anonymous systems can not be implemented (even in synchronous systems). Hence, the paper provides the first consensus algorithm for anonymous systems with this model of partial synchrony and a majority of correct processes.

Keywords: Agreement problem, Asynchrony, Consensus, Distributed computability, Failure detector, Homonymous systems, Message-passing, Process crash

1. Introduction

Homonymous systems Distributed computing is on mastering uncertainty created by adversaries. The first adversary is of course the fact that the processes are geographically distributed which makes impossible to instantaneously obtain

Email addresses: sergio.arevalo@eui.upm.es (Sergio Arévalo), antonio.fernandez@imdea.org (Antonio Fernández Anta), damien.imbs@gmail.com (Damien Imbs), ernes@eui.upm.es (Ernesto Jiménez), michel.raynal@irisa.fr (Michel Raynal)

¹All messages sent after some bounded time GST will be received after at most an unknown bounded latency δ .

a global state of the system. An adversary can be static (e.g., synchrony or anonymity) or dynamic (e.g., asynchrony, mobility, etc.). The net effect of asynchrony and failures is the most studied pair of adversaries.

This paper is on agreement in crash-prone message-passing distributed systems. While this topic has been deeply investigated in the past in the context of asynchrony and process failures (e.g., [18, 20]), we additionally consider here that several processes can have the same identity, i.e., the additional static adversary that is *homonymy*. Systems that work well in the presence of homonymy can be useful in a number of practical situations. For instance, they can tolerate misconfigurations of the processes that result in multiple processes with the same id. Moreover, homonymy can be included in the system design, for instance simply using a default identifier in all processes (which leads to an anonymous system) or using independently randomly generated values as processes id (so that the same id can be chosen by more than one process). In large systems (like peer-to-peer), avoiding the burden of guaranteeing unique identifiers may compensate the cost of dealing with homonymy. One more application example is provided in [14] where users keep their privacy by taking their domain as their identifier (the same identifier is then assigned to all the users of the same domain).

Observe that homonymy is a generalization of two cases: (1) having unique identifiers and (2) having the same identifier for all the processes (anonymity), which are the two extremes of homonymy. However, most algorithms proposed for classical systems with unique identifiers do not work correctly in the presence of id collisions. On the other hand, it has been shown [14, 7] that systems with enough different process identifiers may solve more problems or may have better performance than anonymous systems (where implicitly all identifiers are the same).

We also assume that the distributed system has to face another static adversary, which is the fact that, initially, each process only knows its own identity. This is what is meant by saying that the system has to work *without initial knowledge of the membership*. This static adversary has been recently identified as of significant relevance in certain distributed contexts [17].

How to face adversaries It is well-known that lots of problems cannot be solved in presence of some adversaries (e.g., [1, 2, 15, 21]). When considering process crash failures, the *failure detector* approach introduced in [9, 10] (see [19] for an introductory presentation) has proved to be very attractive. It allows to enrich an otherwise too poor distributed system to solve a given problem P , in order to obtain a more powerful system in which P can be solved.

A failure detector is a distributed oracle that provides processes with information related to failed processes, and can consequently be used to enrich the computability power of asynchronous send/receive message-passing systems. According to the type (set of process identities, integers, etc.) and the quality of this information, several failure detector classes have been proposed. We refer the reader to [20] where classes of failure detectors suited to agreement and communication problems, corresponding failure detector-based algorithms, and additional behavioral assumptions that (when satisfied) allow these failure detectors to be implemented are presented. It is interesting to observe that none of the original failure detectors introduced in [10] can be implemented without initial knowledge of the membership [17].

Aim of the paper Agreement problems are central as soon as one wants to capture the essence of distributed computing. (If processes do not have to agree in one way or another, the problem we have to solve is not a distributed computing problem!) The aim of this paper is consequently to understand the type of information on failures that is needed when one has to solve an agreement problem in presence of asynchrony, process crashes, homonymy, and lack of initial knowledge of the membership. As consensus is the most central agreement problem we focus on it.

Related work As far as we know, consensus in anonymous networks has been addressed first in [5, 13] ([13] considers different synchrony assumptions while [5] considers systems enriched with failure detectors). Connectivity requirements for agreement in anonymous networks is addressed in [16].

To the best of our knowledge, up to now agreement in homonymous systems has been addressed only in [14] and [7]. In the former paper the authors consider that, among the n processes, up to t of them can commit Byzantine failures. The system is homonymous in the sense that there are ℓ , $1 \leq \ell \leq n$, different authenticated identities, each process has one identity, and several processes can share the same identity. It is shown in that paper that $\ell > 3t$ and $\ell > \frac{3t+n}{2}$ are necessary and sufficient conditions for solving consensus in synchronous systems and partially synchronous systems, respectively. The latter paper [7] mainly explores consensus in a shared memory system with anonymous processes, and bounds the complexity (namely, individual write and step complexities) of solving consensus with the aid of an anonymous leader elector $A\Omega$ (see below). They show that if the system is homonymous instead of purely anonymous these bounds can be improved.

The consensus problem in anonymous asynchronous crash-prone message-passing systems has been recently addressed in [5] (for the first time to our knowledge). In such systems, processes have no identity at all¹. This paper introduces an anonymous counterpart² (denoted \overline{AP} later in [6]) of the perfect failure detector P introduced in [10]. A failure detector of class \overline{AP} returns an upper bound (that eventually becomes tight) of the current number of alive processes. The paper then shows that there is an inherent price associated with anonymous consensus, namely, while the lower bound on the number of rounds in a non-anonymous system enriched with P is $t+1$ (where t is the maximum number of faulty processes), it is $2t+1$ in an anonymous system enriched with \overline{AP} . The algorithm proposed assumes knowledge of the parameter t .

More general failure detectors suited to anonymous distributed systems are presented in [6]. Among other results, this paper introduces the anonymous counterpart $A\Sigma$ of the quorum failure detector class Σ [12] and the anonymous counterpart $A\Omega$ of the eventual leader failure detector class Ω [9]. It also presents the failure detector class AP which is the complement of \overline{AP} . An important result of [6] is the fact that relations linking failure detector classes are not the

¹They must also execute the same program, because otherwise they could use the program (or a hash of it) as their identity. We consider that it is the same if processes have no identity or they have the same identity for all processes, since a process that lacks an identity can choose a default value (e.g., \perp) as its identifier.

²In this paper, when we say that a failure detector A is the *counterpart* of a failure detector B we mean that, in a classical asynchronous system (i.e., where each process has its own identity) enriched with a failure detector of class A , it is possible to design an algorithm that builds a failure detector of the class B and vice-versa by exchanging A and B . Said differently, A and B have the same computability power in a classical crash-prone asynchronous system.

same in non-anonymous systems and anonymous systems. This is also the case if processes do not know the number n of processes in the system (unknown membership in anonymous systems). If n is unknown, the equivalence between AP and \overline{AP} , shown in [6], does not hold anymore.

Regarding implementability, it is stated in [6] that $A\Omega$ is not *realistic* (i.e., it can not be implemented in an anonymous synchronous system [11]). If the membership is unknown, it is not hard to show that AP is not realistic either, applying similar techniques as those in [17]. On the other hand, while \overline{AP} can be implemented in an anonymous synchronous system, it is easy to show that it cannot be implemented in most partially synchronous systems (e.g., in particular, in those with all links eventually timely).

Recently, in [8] two new failure detectors for anonymous systems have been proposed, called $A\Omega'$ and $A\Sigma'$, which combined are claimed to be the weakest failure detector to solve consensus in anonymous systems.

Contributions As mentioned, we explore the consensus problem in homonymous systems. Additional adversaries considered are asynchrony, process crashes, and lack of initial knowledge of the membership. We can summarize the main contributions of this paper as follows.

First, the paper defines new classes of failure detectors suited to homonymous systems. These classes, denoted $H\Omega$ and $H\Sigma$, are shown to be homonymous counterparts of Ω and Σ , respectively. The interest on the latter classes is motivated by the fact that $\langle \Sigma, \Omega \rangle$ is the weakest failure detector to solve consensus in crash prone asynchronous message-passing systems for any number of process failures [12]. The paper also investigates the relations linking $H\Sigma$, $A\Sigma$ and Σ , and shows that both $H\Omega$ and $H\Sigma$ can be obtained from \overline{AP} in asynchronous anonymous systems. As a byproduct, we also introduce a new failure detector class denoted $\diamond H\overline{P}$, that is the homonymous counterpart of $\diamond\overline{P}$ (the complement of $\diamond P$ [10]), which we consider of independent interest.

Then, the paper explores the implementability of these classes of failure detectors. It presents an implementation of $\diamond H\overline{P}$ in homonymous message-passing systems with partially synchronous processes and eventually timely links. This algorithm does not require that the processes know the system membership. Since $H\Omega$ can be trivially implemented from $\diamond H\overline{P}$ without communication, $H\Omega$ is realistic and can also be implemented in a partially synchronous homonymous system without membership knowledge. The paper also presents an implementation of $H\Sigma$ in a synchronous homonymous message-passing system without membership knowledge.

Finally, the paper presents two consensus algorithms for asynchronous homonymous systems enriched with $H\Omega$. Both algorithms are derived from consensus algorithms for anonymous systems proposed in [4] and [6], respectively. The main challenge, and hence, the main contribution of our algorithms, is to modify the original algorithms that used $A\Omega$ to use $H\Omega$ instead. In the second algorithm, also the use of $A\Sigma$ has been replaced by the use of $H\Sigma$.

The first algorithm assumes that each process knows the value n and that a majority of processes is correct in all executions³. Since, as mentioned, $H\Omega$ can be implemented with partial synchrony, the combination of the algorithms

³The knowledge of n can be replaced by the knowledge of a parameter α such that, $\alpha > n/2$ and, in all executions, at least α processes are correct.

presented (to implement $H\Omega$ and to solve consensus with $H\Omega$) form a distributed algorithm that solves consensus in any homonymous system with partially synchronous processes, eventually timely links, and a majority of correct processes. When the system is anonymous, this result relaxes the known conditions to solve consensus, since previous algorithms were based on unrealistic failure detectors ($A\Omega$) or failure detectors that require a larger degree of synchrony (\overline{AP}).

The second consensus algorithm presented works for any number of process crashes, and does not need to know n , but assumes that the system is enriched with the pair of failure detectors $\langle H\Sigma, H\Omega \rangle$. This algorithm, combined with the algorithms to implement $H\Sigma$ and $H\Omega$, shows that the consensus problem can be solved in *synchronous* homonymous systems subject to any number of crash failures without the initial knowledge neither of the parameter t nor of the membership. Applied to anonymous systems, this result relaxes the known conditions to solve consensus under any number of failures, since previous algorithms used unrealistic detectors ($A\Omega$) or required to know t or an upper bound on it.

Roadmap The paper is made up of 5 sections. Section 2 presents the system model. Section 3 introduces failure detector classes suited to homonymous systems, and explores their relation with other classes and their implementability. Finally, Section 5 presents failure detector-based homonymous consensus algorithms.

2. System Model

Homonymous processes We consider a distributed system with n processes. Let Π denote the set of processes with $|\Pi| = n$. We use $id(p)$ to denote the identity of process $p \in \Pi$. The system is *homonymous*, which means that different processes may have the same identity. More formally, $p \neq q \not\Rightarrow id(p) \neq id(q)$. Two processes with the same identity are said to be *homonymous*. Let $S \subseteq \Pi$ be any subset of processes. We define $I(S)$ as the *multiset* (sometimes also called *bag*) of process identities in S , $I(S) = \{id(p) : p \in S\}$. Let us remember that, differently from a set, an element of a multiset can appear more than once. Hence, as $I(S)$ may contain several times the same identity, we always have $|I(S)| = |S|$. The *multiplicity* (number of instances) of identity i in a multiset I is denoted $mult_I(i)$. When I is clear from the context we will use simply $mult(i)$. $P(I) \subseteq \Pi$ is used to denote the processes whose identity is in the multiset I , i.e., $P(I) = \{p : p \in \Pi \wedge id(p) \in I\}$. We assume that two homonymous processes execute the same program, because otherwise they could use the program (or a hash of it) as a way to differentiate their identities.

We assume that the system works *without initial knowledge of the membership*. This formally means that an algorithm executed by a process $p \in \Pi$ can only initially use its own identity $id(p)$, and cannot use the identity of any other process nor the system membership $I(\Pi)$, unless learned during its execution (by exchanging messages). Moreover, only in Subsection 5.2 it is assumed that the system size n and an upper bound t on the number of faulty processes are initially known. Observe that the set Π is a formalization tool that is not known by the set of processes of the system.

Processes are asynchronous, unless otherwise stated. We assume that time advances at discrete steps. We assume

a global clock whose values are the positive natural numbers, but processes cannot access it. Processes can fail by crashing, i.e., stop taking steps. A process that crashes in a run is said to be *faulty* and a process that is not faulty in a run is said to be *correct*. The set of correct processes is denoted by $Correct \subseteq \Pi$. A process that has not crashed (yet) at a given time τ is *alive* at that time. Note that a correct process is always alive.

Communication The processes can invoke the primitive $broadcast(m)$ to send a message m to all processes of the system (including itself). This communication primitive is modeled in the following way. The network is assumed to have a directed link from process p to process q for each pair of processes $p, q \in \Pi$ (p does not need to be different from q). Then, $broadcast(m)$ invoked at process p sends one copy of message m along the link from p to q , for each $q \in \Pi$. The receiving process q cannot identify the link through which a message was received.

Unless otherwise stated, links are asynchronous and reliable, i.e., links neither lose messages nor duplicate messages nor corrupt messages nor generate spurious messages. If a process crashes while broadcasting a message, the message is received by an arbitrary subset of processes.

Notation and time-related definitions The previous model is denoted $HAS[\emptyset]$ (Homonymous Asynchronous System). We use $HPS[\emptyset]$ to denote a homonymous system where processes are partially synchronous and links are eventually timely. A process is *partially synchronous* if the time to execute a step is bounded, but the bound is unknown. A link is *eventually timely* if there is an unknown global stabilization time (denoted GST) after which all messages sent across the link are delivered in a bounded δ time, where δ is unknown. Messages sent before GST can be lost or delivered after an arbitrary (but finite) time. Finally, we use $HSS[\emptyset]$ to denote a Homonymous Synchronous System, which is one in which links are reliable, there are bounds on the time to execute a step and the latency of every message, and these bounds are known.

$AS[\emptyset]$ denotes the classical asynchronous system with unique identities and reliable channels. Finally, $AAS[\emptyset]$ denotes the Anonymous Asynchronous System model [6]. Observe that $AS[\emptyset]$ and $AAS[\emptyset]$ are special cases (actually extreme cases with respect to homonymy) of $HAS[\emptyset]$ (an anonymous system can be seen as a homonymous system where all processes have the same default identifier \perp).

3. Failure Detectors

In this section we define failure detectors previously proposed and the ones proposed here for homonymous systems. Then, relationships between these detectors are derived, and their implementability is explored.

3.1. Failure detectors for classical and anonymous systems

We briefly describe here some failure detector previously proposed. We start with the classes that have been defined for $AS[\emptyset]$. (Observe that a variable X that is local to a process p is marked with p as subindex, X_p . If the variable has a superindex τ , X^τ , this indicates that we consider the variable at time τ .)

A failure detector of class Σ [12] provides each process $p \in \Pi$ with a variable $trusted_p$ which contains a multiset⁴ of process identifiers. The properties that are satisfied by these multisets are [Liveness] $\forall p \in Correct, \exists \tau \in \mathbb{N} : \forall \tau' \geq \tau, trusted_p^{\tau'} \subseteq I(Correct)$, and [Safety] $\forall p, q \in \Pi, \forall \tau, \tau' \in \mathbb{N}, trusted_p^\tau \cap trusted_q^{\tau'} \neq \emptyset$.

A failure detector of class Ω [9] provides each process $p \in \Pi$ with a variable $leader_p$ such that [Election] eventually all these variables contain the same process identifier of a correct process.

The following failure detector classes have been defined for anonymous systems $AAS[\emptyset]$.

A failure detector of class $A\Omega$ [6] provides each process $p \in \Pi$ with a variable a_leader_p , such that [Election] there is a time after which, permanently, (1) there is a correct process whose Boolean variable is true, and (2) the Boolean variables of the other correct processes are false.

A failure detector of class \overline{AP} [5] provides each process $p \in \Pi$ with a variable $anap_p$ such that, if $anap_p^\tau$ and $Correct^\tau$ denote the value of this variable and the number of alive processes at time τ , respectively, then [Safety] $\forall p \in \Pi, \forall \tau \in \mathbb{N}, anap_p^\tau \geq |Correct^\tau|$, and [Liveness] $\exists \tau \in \mathbb{N}, \forall p \in Correct, \forall \tau' \geq \tau, anap_p^{\tau'} = |Correct|$.

A failure detector of class $A\Sigma$ [6] provides each process $p \in \Pi$ with a variable a_sigma_p that contains a set of pairs of the form (x, y) . The parameter x is a label provided by the failure detector, and y is an integer. Let us denote $a_sigma_p^\tau$ the value of variable a_sigma_p at time τ . Let $S_A(x) = \{p \in \Pi \mid \exists \tau \in \mathbb{N} : (x, -) \in a_sigma_p^\tau\}$. Any failure detector of class $A\Sigma$ must satisfy the following properties:

- Validity. No set a_sigma_p ever contains simultaneously two pairs with the same label.
- Monotonicity. $\forall p \in \Pi, \forall \tau \in \mathbb{N} : ((x, y) \in a_sigma_p^\tau) \implies (\forall \tau' \geq \tau : \exists y' \leq y : (x, y') \in a_sigma_p^{\tau'})$.
- Liveness. $\forall p \in Correct, \exists \tau \in \mathbb{N} : \forall \tau' \geq \tau : \exists (x, y) \in a_sigma_p^{\tau'} : (|S_A(x) \cap Correct| \geq y)$.
- Safety. $\forall p_1, p_2 \in \Pi, \forall \tau_1, \tau_2 \in \mathbb{N}, \forall (x_1, y_1) \in a_sigma_{p_1}^{\tau_1} : \forall (x_2, y_2) \in a_sigma_{p_2}^{\tau_2} : \forall T_1 \subseteq S_A(x_1) : \forall T_2 \subseteq S_A(x_2) : ((|T_1| = y_1) \wedge (|T_2| = y_2)) \implies (T_1 \cap T_2 \neq \emptyset)$.

3.2. Failure detectors for homonymous systems

Classical failures detectors [10] output a set of processes' identifiers. Our failures detectors extend this output to a multiset of processes' identifiers, due to the homonymy nature of the system. The following are the new failure detectors proposed for homonymous systems.

A failure detector of class $\diamond H\overline{P}$ eventually outputs forever the multiset with the identifiers of the correct processes. More formally, a failure detector of class $\diamond H\overline{P}$ provides each process $p \in \Pi$ with a variable $h_trusted_p$, such that [Liveness] $\forall p \in Correct, \exists \tau \in \mathbb{N} : \forall \tau' \geq \tau, h_trusted_p^{\tau'} = I(Correct)$. This failure detector $\diamond H\overline{P}$ is the counterpart of $\diamond \overline{P}$.

⁴Note that Σ was originally defined for systems without homonymy, where $trusted_p$ is a set. In a homonymous system, the natural generalization is that $trusted_p$ is a multiset.

A failure detector of class $H\Omega$ eventually outputs the same identifier ℓ and number c at all processes, such that ℓ is the identifier of some correct process, and c is the number of correct processes that have this identifier ℓ . More formally, a failure detector of class $H\Omega$ provides each process $p \in \Pi$ with two variables h_leader_p and $h_multiplicity_p$, such that [Election] $\exists \ell \in I(\text{Correct}), \exists \tau \in \mathbb{N} : \forall \tau' \geq \tau, \forall p \in \text{Correct}, h_leader_p^{\tau'} = \ell$, and $h_multiplicity_p^{\tau'} = \text{mult}_{I(\text{Correct})}(\ell)$.

Any correct process p such that $id(p) = \ell$ is called a *leader*. Note that this failure detector does not choose only one leader, like in Ω or in $A\Omega$, but a set of leaders with the same identifier. When all identifiers are different, the class $H\Omega$ is equivalent to Ω (i.e., any detector in $H\Omega$ can be trivially transformed into a detector in Ω , and vice versa). Furthermore, we have the following observation.

Observation 1. A failure detector of class $H\Omega$ can be obtained from any detector D of class $\diamond H\bar{P}$ without any communication.

The transformation from $\diamond H\bar{P}$ to $H\Omega$ can be done by, for instance, setting at each process p periodically h_leader_p to the smallest element in $D.h_trusted_p$, and $h_multiplicity_p \leftarrow \text{mult}_{D.h_trusted_p}(h_leader_p)$.

A failure detector of class $H\Sigma$ provides each process $p \in \Pi$ with two variables h_quora_p and h_labels_p , where h_quora_p is a set of pairs of the form (x, m) (x is a label, and m is a multiset such that $m \subseteq I(\Pi)$) and h_labels_p is a set of labels. Roughly speaking, each pair (x, m) determines a set of quora, and the set h_labels_p of a process p determines in which of these sets it participates. More formal, let us denote $h_quora_p^\tau$ and $h_labels_p^\tau$ the values of variables h_quora_p and h_labels_p at time τ , respectively. Let $S(x) = \{p \in \Pi \mid \exists \tau \in \mathbb{N} : x \in h_labels_p^\tau\}$. Any failure detector of class $H\Sigma$ must satisfy the following properties:

- **Validity.** No set h_quora_p ever contains simultaneously two pairs with the same label.
- **Monotonicity.** $\forall p \in \Pi, \forall \tau \in \mathbb{N}, \forall \tau' \geq \tau$: (1) $h_labels_p^\tau \subseteq h_labels_p^{\tau'}$, and (2) $((x, m) \in h_quora_p^\tau) \implies \exists m' \subseteq m : (x, m') \in h_quora_p^{\tau'}$.
- **Liveness.** $\forall p \in \text{Correct}, \exists \tau \in \mathbb{N} : \forall \tau' \geq \tau, \exists (x, m) \in h_quora_p^{\tau'} : m \subseteq I(S(x) \cap \text{Correct})$.
- **Safety.** $\forall p_1, p_2 \in \Pi, \forall \tau_1, \tau_2 \in \mathbb{N}, \forall (x_1, m_1) \in h_quora_{p_1}^{\tau_1} : \forall (x_2, m_2) \in h_quora_{p_2}^{\tau_2} : \forall Q_1 \subseteq S(x_1), \forall Q_2 \subseteq S(x_2), (I(Q_1) = m_1 \wedge I(Q_2) = m_2) \implies (Q_1 \cap Q_2 \neq \emptyset)$.

Example of $H\Sigma$. For instance, consider a system with process set $\Pi = \{1, 2, 3\}$, whose identifiers are $id(1) = A$, $id(2) = A$, and $id(3) = B$. (Observe that $I(\Pi) = \{A, A, B\}$.) Assume that at time τ the processes have $h_labels_1^\tau = \{la, lc\}$, $h_labels_2^\tau = \{la, lb\}$, and $h_labels_3^\tau = \{lb, lc\}$, where la , lb , and lc are labels. If these variables do not change after time τ , they determine the following quora $S(la) = \{1, 2\}$, $S(lb) = \{2, 3\}$, $S(lc) = \{1, 3\}$.


```

1 Init
2  $h\_labels_p \leftarrow \{s : (s \subseteq I(\Pi)) \wedge (id(p) \in s)\};$ 
3  $h\_quora_p \leftarrow \emptyset;$ 
4 repeat forever
5    $q \leftarrow D.trusted_p;$ 
6    $h\_quora_p \leftarrow h\_quora_p \cup \{(q, q)\}$ 
7 end repeat.

```

Figure 1: Algorithm to transform $D \in \Sigma$ to $H\Sigma$ with initial knowledge of membership (code for process p).

Assume that process 2 is faulty. The liveness property must ensure that processes 1 and 3 eventually have appropriate values in variables h_quora . Since $I(S(la) \cap Correct) = A^5$, $I(S(lb) \cap Correct) = B$, and $I(S(lc) \cap Correct) = AB$, all the pairs (la, A) , (lb, B) , (lc, AB) , (lc, A) , and (lc, B) satisfy the predicate of the liveness property defined over pairs (x, m) . Hence, if for instance $h_quora_1 = \{(lb, B)\}$ and $h_quora_3 = \{(la, AB), (lc, AB)\}$, the liveness property is satisfied. (Note that the monotonicity property guarantees that (lb, B) will always be in h_quora_1 and that one of (lc, AB) , (lc, A) or (lc, B) will always be in h_quora_3 .)

The safety property also holds between these two particular instances of h_quora_1 and h_quora_3 . For pair (lb, B) the only $Q_1 \subseteq S(lb)$ such that $I(Q_1) = B$ is set $Q_1 = \{3\}$. For pair (la, AB) there is no $Q_2 \subseteq S(la)$ such that $I(Q_2) = AB$. For pair (lc, AB) the only appropriate Q_2 is set $Q_2 = S(lc) = \{1, 3\}$. Obviously $Q_1 \cap Q_2 \neq \emptyset$. Note that this is necessary but not sufficient for the safety property to hold, since all the values of variables h_quora_1 , h_quora_2 , and h_quora_3 at all times must be considered.

Comparing $H\Sigma$ and $A\Sigma$, one can observe that $H\Sigma$ has pairs (x, m) in which m is a multiset of identifiers, while $A\Sigma$ uses pairs (x, y) in which y is an integer. However, a more important difference is that, in $H\Sigma$, each process has two variables. Then, the labels that a process p has in h_quora_p can be disconnected from those it has in h_labels_p . For instance, in the above example, the pair (lb, B) is in h_quora_1 while process 1 is not in the quorum $S(lb)$. This is not possible in $A\Sigma$, where a process that has pair (x, y) always belongs to the quorum identified by x .

3.3. Reductions between failure detectors

In this section we claim that it can be shown, via reductions, the relation of the newly defined failure detector classes with the previously defined classes. We use the standard form of comparing the relative power of failure detector classes of [10]. A failure detector class X is *stronger* than class X' in system $Y[\emptyset]$ if there is an algorithm A that emulates the output of a failure detector of class X' in $Y[X]$ (i.e., system $Y[\emptyset]$ enhanced with a failure detector D of class X). We also say that X' can be obtained from X in $Y[\emptyset]$. Two classes are equivalent if this property can be shown in both directions.

3.3.1. From Σ to $H\Sigma$

We prove that, if identifiers are unique, a detector of class $H\Sigma$ can be obtained from any detector D of class Σ .

⁵We omit the brackets in multisets to simplify the notation.

```

1 Init
2  $h\_labels_p \leftarrow \emptyset$ ;
3  $h\_quora_p \leftarrow \emptyset$ ;
4  $mship_p \leftarrow \emptyset$ ;
5 start tasks T1 and T2;
6 Task T1
7 repeat forever
8   broadcast ( $IDENT, id(p)$ );
9    $q \leftarrow D.trusted_p$ ;
10   $h\_quora_p \leftarrow h\_quora_p \cup \{(q, q)\}$ 
11 end repeat.
12
13 Task T2
14 upon reception of ( $IDENT, i$ ) do
15    $mship_p \leftarrow mship_p \cup \{i\}$ ;
16    $h\_labels_p \leftarrow \{s : (s \subseteq mship_p) \wedge (id(p) \in s)\}$ .

```

Figure 2: Algorithm to transform $D \in \Sigma$ to $H\Sigma$ without initial knowledge of membership (code for process p).

Theorem 1. *A failure detector of class $H\Sigma$ can be obtained from any detector D of class Σ in a system with unique identifiers, under either one of the following conditions:*

1. *without any communication if every process initially knows the membership $I(\Pi)$, or*
2. *in system $AS[\Sigma]$ (the membership does not need to be known initially).*

Proof: Let $D.trusted_p$ be the variable of Σ failure detector D at process p . Figures 1 and 2 present the algorithms to transform D into a failure detector of class $H\Sigma$ in Cases 1 and 2, respectively. In both cases, at each process p initially $h_quora_p \leftarrow \emptyset$, and infinitely often this variable is updated with the following sentences: $q \leftarrow D.trusted_p$, and $h_quora_p \leftarrow h_quora_p \cup \{(q, q)\}$. In Case 1, initially every process p sets $h_labels_p \leftarrow \{s : (s \subseteq I(\Pi)) \wedge (id(p) \in s)\}$ and it never changes it in the run. In Case 2, every process p initially sets $h_labels_p \leftarrow \emptyset$, and repeatedly broadcasts a message $IDENT(id(p))$. Process p also has a variable $mship_p$ initially set to $mship_p \leftarrow \emptyset$. After receiving a message $IDENT(i)$, process p updates $mship_p \leftarrow mship_p \cup \{i\}$, and $h_labels_p \leftarrow \{s : (s \subseteq mship_p) \wedge (id(p) \in s)\}$.

We prove now the properties of $H\Sigma$:

- **Validity.** Since h_quora_p is a set, and the elements included in it are of the form (q, q) (see Line 5 in Figure 1, and Line 10 in Figure 2) there can not be two pairs with the same label.
- **Monotonicity.** The monotonicity of h_labels_p in Figure 1 is obvious because it is initialized in Line 2 and never changes. With respect to Figure 2, h_labels_p is initially empty, and it is related with the set $mship_p$, such that if $mship_p$ grows then h_labels_p either grows or remains the same. Hence h_labels_p never decreases because $mship_p$ never decreases (see Line 15 in Figure 2). The monotonicity of h_quora_p in Figures 1 and 2 follows from the fact that h_quora_p is initially empty, and any element (q, q) included in it is never removed.
- **Liveness.** Consider any correct process p . In Figure 2, eventually, $Correct \subseteq mship_p$ permanently (from the exchange of $IDENT$ messages and Line 15 of Figure 2). Then, in both algorithms eventually $\{s : (s \subseteq$

```

1 Init
2  $alive_p \leftarrow$  empty list;
3 start Tasks T1 and T2;
4 Task T1
5 repeat forever
6   broadcast ( $ALIVE, id(p)$ )
7 end repeat.
8
9 Task T2
10 upon reception of ( $ALIVE, i$ ) do
11   if  $i \in alive_p$  then move  $i$  to the first position of  $alive_p$ 
12   else insert  $i$  in the first position of  $alive_p$ 
13   end if.

```

Figure 3: Algorithm to implement a failure detector of class Ξ without initial knowledge of membership in $AS[\emptyset]$ (code for process p).

$I(Correct) \wedge (id(p) \in s) \subseteq h_labels_p$ permanently (from Line 2 in Figure 1, and Line 16 in Figure 2). Hence, there is a time τ after which, for every set $s \subseteq I(Correct)$, $I(S(s)) = s$ and $S(s) \subseteq Correct$.

The Liveness property of Σ guarantees that, at some time $\tau' \geq \tau$, the variable q is assigned a set s that contains only correct processes and (s, s) will be included in h_quora_p after that. Therefore, there is a time after which h_quora_p contains (s, s) permanently (from monotonicity). Since $s \subseteq I(S(s) \cap Correct) = I(S(s)) = s$, the property follows.

- **Safety.** Consider two pairs $(x_1, m_1) \in h_quora_{p_1}^{\tau_1}$ and $(x_2, m_2) \in h_quora_{p_2}^{\tau_2}$, for any $p_1, p_2 \in \Pi$ and any $\tau_1, \tau_2 \in \mathbb{N}$. From the management of the h_quora variables (Lines 3, 5, and 6 in Figure 1, and Lines 3, 9, and 10 in Figure 2), we have that m_1 and m_2 are values taken from $D.trusted_{p_1}$ and $D.trusted_{p_2}$, respectively. Hence, the sets m_1 and m_2 must intersect from the Safety property of the Σ failure detector D . Then, if $I(Q_1) = m_1$ and $I(Q_2) = m_2$, given that we are in a system with unique identifiers, Q_1 and Q_2 must intersect. ■

3.3.2. From $H\Sigma$ to Σ

We define now a new class of failure detector that will be used for reductions between the above failure detector classes. While the service provided by this detector has been already used [22, 6], it was never formally defined. The new failure detector class, denoted Ξ , will only be defined for systems with unique identifiers, i.e., non homonymous.

Definition 1. A failure detector of class Ξ provides each process $p \in \Pi$, in a system with unique process identifiers, with a variable $alive_p$ which contains a sequence of process identifiers. Any failure detector of class Ξ must satisfy the following property:

- **Liveness.** Eventually, the identifiers of the correct processes are permanently in the prefix of $alive_p$. More formally, let $rank(i, alive_p^{\tau})$ denote the position (starting from 1) of process identifier i in $alive_p^{\tau}$ (with $rank(i, alive_p^{\tau}) =$

```

1 Init
2 start Tasks T1 and T2;
3 Task T1
4 repeat forever
5   broadcast (LABELS, id(p), D.h.labelsp);
6   if  $\exists(x, m) \in D.h\_quora_p : (idents_p[x] \text{ has been created}) \wedge (m \subseteq idents_p[x])$  then
7     let candidatesp = {m : ((x, m)  $\in D.h\_quora_p$ )  $\wedge$  (identsp[x] has been created)  $\wedge$  (m  $\subseteq idents_p[x]$ )};
8     trustedp  $\leftarrow$  any m  $\in candidates_p$  with smallest  $\max_{i \in m} rank(i, X.alive_p)$ 
9   end if
10 end repeat.
11
12 Task T2
13 upon reception of (LABELS, i,  $\ell$ ) do
14   foreach  $x \in \ell$  do
15     if identsp[x] has not been created then create identsp[x]  $\leftarrow \emptyset$  end if;
16     identsp[x]  $\leftarrow idents_p[x] \cup \{i\}$ 
17   end foreach.

```

Figure 4: Algorithm to transform $D \in H\Sigma$ to Σ in a system with unique identifiers, but without initial knowledge of membership (code for process p). The algorithm uses a failure detector X of class Ξ .

∞ if $i \notin alive_p^\tau$). Then, $\forall p \in Correct$, $\exists \tau \in \mathbb{N} : \forall \tau' \geq \tau$, $\forall q \in Correct$, $rank(id(q), alive_p^{\tau'}) \leq |Correct|$.

Observe that the position of the same identifier can be different at different processes, and can vary over time in the same process. From the algorithm of Figure 3, we obtain the following lemma.

Lemma 1. *A failure detector of class Ξ can be implemented in $AS[\emptyset]$ (an asynchronous system with unique identifiers), even when the membership is not known initially.*

Proof: For each process $q \in Correct$, eventually some message $ALIVE(id(q))$ will be received at each process $p \in Correct$. Then $id(q)$ will be included in $alive_p$ and never removed after that. Given any faulty process r , p will stop receiving messages from r by some time τ . Then, after τ process p will never receive a message $ALIVE(id(r))$ and $id(r)$ will never be moved to (inserted in) the first position of $alive_p$. However, after τ , eventually p will receive messages $ALIVE(id(q))$ from each process $q \in Correct$, and each identifier $id(q)$ will be moved to (or inserted in) the first position of $alive_p$. Then, there is some time $\tau' > \tau$ such that, at all times $\tau'' > \tau'$, $rank(id(q), alive_p^{\tau''}) < rank(id(r), alive_p^{\tau''})$. Since this holds for all $p, q \in Correct$ and all $r \notin Correct$, the claim follows. ■

We now show, using the algorithm of Figure 4, that Σ can be obtained from $H\Sigma$ without initial knowledge of the membership. The logic of the algorithm of Figure 4 is somewhat similar to that of the algorithm in Figure 2 in [6].

Theorem 2. *A failure detector of class Σ can be obtained from any detector D of class $H\Sigma$ in $AS[H\Sigma]$ (an asynchronous system with unique identifiers), even when the membership is not known initially.*

Proof: From Lemma 1, we can have a failure detector of class Ξ in an asynchronous system. The condition in Line 6 guarantees that the variable $trusted_p$ is assigned a set of identifiers m only if (x, m) is in h_quora_p , and every process

q whose identifier is in m has x in its set h_labels_q (from the management of the sets $idents_p$). Combining this condition with the safety property of $H\Sigma$ we guarantee the safety property of Σ . The liveness property of Σ holds from the liveness property of $H\Sigma$, the choice of m done in Line 8, and the properties of the failure detector class Ξ as follows. If $p \in Correct$, from the liveness of $H\Sigma$, eventually every time Line 8 is executed, there is some $m \in candidates_p$ with only correct processes. If the failure detector X of class Ξ has already all the correct processes in the lowest ranks of $X.alive_p$ (which eventually happens from its liveness property), then any set m in $candidates_p$, whose largest rank in $X.alive_p$ is minimal, contains only correct processes (which yields the liveness of Σ). ■

The following result shows that, in classical systems with unique identifiers, Σ , $H\Sigma$, and $A\Sigma$ are equivalent.

Corollary 1. *Failure detector classes Σ , $H\Sigma$, and $A\Sigma$ are equivalent in $AS[\emptyset]$.*

Proof: From Theorems 1 and 2 we have that Σ and $H\Sigma$ are equivalent. The equivalence between Σ and $A\Sigma$ was shown in [6]. ■

3.3.3. From $A\Sigma$ to $H\Sigma$

In anonymous systems we have the following properties. Recall that an anonymous system is assumed to be a homonymous system in which every process has a default identifier \perp ⁶.

We show now how to obtain a failure detector of class $H\Sigma$ from a detector of class $A\Sigma$.

Theorem 3. *Class $H\Sigma$ can be obtained from class $A\Sigma$ in $AAS[\emptyset]$ without communication.*

Proof: Let D be a detector of class $A\Sigma$. The transformation can be done as follows. Let \perp be the “default” identifier. Let us denote with \perp^r a multiset of r identifiers \perp . Each process p periodically does as follows. For each pair $(x, y) \in D.a_sigma_p$, the label x is included in h_labels_p and the pair (x, \perp^y) is included in h_quora_p (replacing any pair $(x, -)$ that h_quora_p may contain). The properties of $H\Sigma$ follow trivially from the properties of $A\Sigma$. ■

3.3.4. From \overline{AP} to $\diamond H\overline{P}$ and $H\Sigma$

We show here how failure detectors of the classes $\diamond H\overline{P}$ and $H\Sigma$ can be obtained for a failure detector of class \overline{AP} without communication.

Lemma 2. *A failure detector of class $\diamond H\overline{P}$ can be obtained from any detector D of class \overline{AP} in $AAS[\emptyset]$ (an anonymous asynchronous system) without communication.*

⁶Note that this differs from the assumption used in [6].

Proof: The transformation can be done as follows. Let \perp be the “default” identifier. Each process p periodically updates $h_trusted_p$ to a multiset of $D.anap_p$ identifiers \perp . The liveness property of D guarantees the liveness property of $\diamond H\bar{P}$. ■

Lemma 3. *A failure detector of class $H\Sigma$ can be obtained from any detector D of class \bar{AP} in $AAS[\emptyset]$ (an anonymous asynchronous system) without communication.*

Proof: The transformation can be done as follows. Let \perp be the “default” identifier. Let us denote with \perp^r a multiset of r identifiers \perp . Each process p periodically does as follows. After obtaining a value y from $D.anap_p$, the label \perp^y is included in h_labels_p and the pair (\perp^y, \perp^y) is included in h_quora_p . The Validity and Monotonicity of $H\Sigma$ hold trivially. Liveness follows since, from the safety of \bar{AP} , only correct processes see an output of $D.anap = c = |Correct|$, and from the liveness property all of them do it. Then, every correct process p eventually inserts \perp^c in h_labels_p and (\perp^c, \perp^c) in h_quora_p , and only those processes. Safety of $H\Sigma$ comes from the safety property of \bar{AP} : if, for any y and y' with $y \geq y'$, $|S(\perp^y)| = y$ and $|S(\perp^{y'})| = y'$ (none can be larger), then $S(\perp^y) \subseteq S(\perp^{y'})$. ■

Theorem 4. *Classes $\diamond H\bar{P}$ and $H\Sigma$ can be obtained from class \bar{AP} in $AAS[\emptyset]$ without communication.*

Proof: The proof of Theorem 4 follows from the two previous lemmas 2 and 3. ■

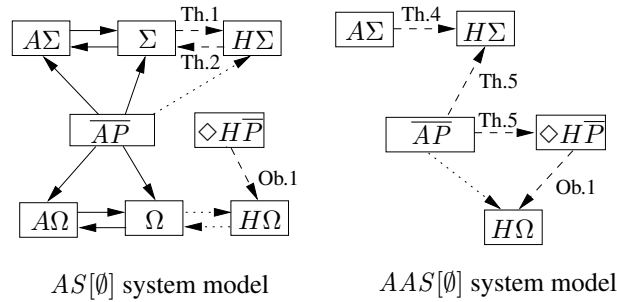


Figure 5: Relations between failure detector classes in the models $AS[\emptyset]$ and $AAS[\emptyset]$. There is an arrow from class X to X' if X is stronger than X' . Solid arrows are relations shown by Bonnet and Raynal in [6]. Dotted arrows are trivial relations. Dashed arrows are relations shown here (the arrow label shows the theorem or observation where the relation is proven).

4. Implementing Failure Detectors in Homonymous Systems

In this section, we show that there are algorithms that implement the failure detectors classes $\diamond H\bar{P}$ and $H\Omega$ in $HPS[\emptyset]$ (homonymous partially synchronous system). We also implement the failure detector $H\Sigma$ in $HSS[\emptyset]$ (homonymous synchronous system). In all cases they do not need to know initially the membership.

```

1 Init
2  $h\_trusted_p \leftarrow \emptyset$ ; // multiset of process identifiers
3  $mship_p \leftarrow \emptyset$ ; // set of process identifiers
4  $r_p \leftarrow 1$ ;
5  $timeout_p \leftarrow 1$ ;
6 start Tasks T1 and T2;
7
8 Task T1
9 repeat forever
10 broadcast (POLLING,  $r_p$ ,  $id(p)$ );
11 wait  $timeout_p$  time;
12  $tmp_p \leftarrow \emptyset$ ; //  $tmp_p$  is an auxiliary multiset
13 for each (P_REPLY,  $r$ ,  $r'$ ,  $id(p)$ ,  $id(q)$ ) received
14     with ( $r \leq r_p \leq r'$ ) do
15     add one instance of  $id(q)$  to  $tmp_p$ 
16 end for;
17  $h\_trusted_p \leftarrow tmp_p$ ;
18  $r_p \leftarrow r_p + 1$ 
19 end repeat.
20
21 Task T2
22 upon reception of (POLLING,  $r_q$ ,  $id(q)$ ) do
23 if  $id(q) \notin mship_p$  then
24      $mship_p \leftarrow mship_p \cup \{id(q)\}$ ;
25     create  $latest\_r_p[id(q)]$ ;
26      $latest\_r_p[id(q)] \leftarrow 0$ 
27 end if;
28 if  $latest\_r_p[id(q)] < r_q$  then
29     broadcast(P_REPLY,  $latest\_r_p[id(q)] + 1$ ,  $r_q$ ,  $id(q)$ ,  $id(p)$ )
30 end if;
31  $latest\_r_p[id(q)] \leftarrow \max(latest\_r_p[id(q)], r_q)$ .
32
33 upon reception of (P_REPLY,  $r$ ,  $r'$ ,  $id(p)$ ,  $-$ ) with ( $r < r_p$ ) do
34      $timeout_p \leftarrow timeout_p + 1$ .

```

Figure 6: Algorithm that implements $\diamond HP$ (code for process p).

4.1. Implementation of $\diamond H\bar{P}$ and $H\Omega$

The algorithm of Figure 6 implements $\diamond H\bar{P}$ (and $H\Omega$ with trivial changes) in $HPS[\emptyset]$ where processes are partially synchronous, links are eventually timely, and membership is not known.

Brief description of the algorithm: It is a polling-based algorithm that executes in rounds. At every round r , the Task 1 of each process p broadcasts $(POLLING, r, id(p))$ messages. After a time $timeout_p$, it gathers in the variable tmp_p (and, hence, also in $h_trusted_p$) a multiset with the senders' identifiers id_s of processes from $(P_REPLY, r', r'', id(p), id_s)$ messages received with $r' \leq r \leq r''$.

Task 2 is related with the reception of $POLLING$ and P_REPLY messages. When a process p receives a $(POLLING, r, id(q))$ message from process q , process p has to respond with as many P_REPLY as process q needs to receive up to round r , and not previously sent by process p (Lines 28-30). Note that the P_REPLY messages are piggybacked in only one message (Line 29). Also note that is in variable $latest_r_p[id(q)]$ where p holds the latest round broadcast to $id(q)$. If it is the first time that process p receives a $(POLLING, -, id)$ message from a process with identifier id , then variable $latest_r_p[id]$ is created and initialized to zero (Lines 23-27).

It is important to remark that, for each different identifier id , only one $(P_REPLY, -, -, id(q), id)$ message is broadcast by each process q . So, if processes v and w with $id(v) = id(w) = x$ broadcast two $(POLLING, r, x)$ messages, then each process p only broadcast one (P_REPLY, r', r'', x, q) message with $r' \leq r \leq r''$. Note that eventually (at least after GST time) each P_REPLY message sent by any process has to be received by all correct processes. Hence, eventually processes v and w will receive all P_REPLY messages generated due to $POLLING$ messages.

Finally, Lines 33-34 of Task 2 allow process p to adapt the variable $timeout_p$ to the communication latency and process speed. When process p receives an outdated $(P_REPLY, r, -, id(p), -)$ message (i.e., a message with round r less than current round r_p), then it increases its variable $timeout_p$.

Lemma 4. *Given processes $p \in Correct$ and $q \notin Correct$, there is a round r such that p does not receive any $(P_REPLY, \rho, \rho', id(p), id(q))$ message from q with $\rho' \geq r$.*

Proof: There is a time τ at which q stops taking steps. If q ever sent a $(P_REPLY, -, -, id(p), id(q))$ message, consider the largest x such that q sent message $(P_REPLY, -, x, id(p), id(q))$. Otherwise, let $x = 0$. Then, the claim holds for $r = x + 1$. ■

Lemma 5. *Given processes $p, q \in Correct$, there is a round r such that, for all rounds $r' \geq r$, when p executes the loop of Lines 14-16 with $r_p = r'$, it has received a message $(P_REPLY, \rho, \rho', id(p), id(q))$ from q with $\rho \leq r' \leq \rho'$.*

Proof: Observe that, since p is correct, it will repeat forever the loop of Lines 9-19, with the value of r_p increasing in one unit at each iteration. Hence, p will be sending forever messages $(POLLING, -, id(p))$ after GST with increasing round numbers, that will eventually be received by q . Then, q eventually will send infinite $(P_REPLY, -, -,$

$id(p), id(q)$) messages after GST , with increasing round numbers. Let $(P_REPLY, x, -, id(p), id(q))$ be the first such message sent by q after GST . Then, for each round number $y \geq x$, there is some message $(P_REPLY, \rho, \rho', id(p), id(q))$ sent by q with $\rho \leq y \leq \rho'$, and these messages are delivered at p at most δ time after being sent.

Now, assume for contradiction that for each round $y \geq x$, there is a round $y' \geq y$ such that, when p executes the loop of Lines 14-16 with $r_p = y'$, it has not received the message $(P_REPLY, \rho, \rho', id(p), id(q))$ from q with $\rho \leq y' \leq \rho'$. But, every time this happens, when the message is finally received, r_p has been incremented in Line 18 and, hence, $timeout_p$ is incremented (in Lines 33-34). Then, eventually, by some round r , the value of $timeout_p$ will be greater than $2\delta + \gamma$, where γ is the maximum time that q takes to execute Lines 22-31. Then, p will receive message $(P_REPLY, \rho, \rho', id(p), id(q))$ with $\rho \leq r' \leq \rho'$ before executing the loop of Lines 14-16 with $r_p = r'$, for all $r' \geq r$. We have reached a contradiction and the claim of the lemma follows. ■

Theorem 5. *The algorithm of Figure 6 implements a failure detector of the class $\diamond H\bar{P}$ in a system $HPS[\emptyset]$ (homonymous system where processes are partially synchronous and links are eventually timely), even if the membership is not known initially.*

Proof: Consider a correct process p . From Lemma 4, there is a round r such that p does not receive any $(P_REPLY, \rho, \rho', -, -)$ message with $\rho' \geq r$ from any faulty process. From Lemma 5, there is a round r' such that for all rounds $r'' \geq r'$, when p executes the loop of Lines 14-16 with $r_p = r''$, it has received a $(P_REPLY, \rho, \rho', -, -)$ message with $\rho \leq r'' \leq \rho'$ from each correct process. Hence, for every round $r'' \geq \max(r, r')$ when the Line 17 is executed with $r_p = r''$, the variable $h_trusted_p$ is updated with the multiset $I(Correct)$. ■

We can obtain $H\Omega$ from the algorithm of Fig. 6 without additional communication. This can be done by simply including, immediately after Line 17, $h_leader_p \leftarrow \min(h_trusted_p)$ (i.e., the smallest identifier in $h_trusted_p$) and $h_multiplicity_p \leftarrow mult_{h_trusted_p}(h_leader_p)$.

Corollary 2. *The algorithm of Figure 6 can be changed to implement a failure detector of the class $H\Omega$ in a system $HPS[\emptyset]$ (homonymous system where processes are partially synchronous and links are eventually timely), even if the membership is not known initially.*

4.2. Implementation of $H\Sigma$

Figure 7 implements $H\Sigma$ in $HSS[\emptyset]$ where processes are synchronous, links are timely, and membership is not known.

Brief description of the algorithm It runs in synchronous steps. In each step every process p broadcasts a $(IDENT, id(p))$ message. Then, process p waits for $(IDENT, -)$ messages sent through reliable links in this synchronous step by alive processes. Process p gathers in the multiset variable $mset_p$ the identifiers id of all $(IDENT, id)$ messages

```

1  $h\_labels_p \leftarrow \emptyset$ ;
2  $h\_quora_p \leftarrow \emptyset$ ;
3 for each synchronous step do
4   broadcast ( $IDENT, id(p)$ );
5   wait for the messages sent in this synchronous step;
6    $mset_p \leftarrow$  multiset of identifiers received in ( $IDENT, -$ ) msgs;
7    $h\_quora_p \leftarrow h\_quora_p \cup \{(mset_p, mset_p)\}$ ;
8    $h\_labels_p \leftarrow h\_labels_p \cup \{mset_p\}$ 
9 end for.

```

Figure 7: Algorithm to implement $H\Sigma$ without knowledge of membership (code for process p)

received. At the end of this step, variables h_quora_p and h_labels_p are updated with the value of $mset_p$. Note that for process p the label x of a quorum (x, m) is formed by the multiset $mset_p$ (i.e., $x = m = mset_p$).

Theorem 6. *The algorithm of Figure 7 implements a failure detector of the class $H\Sigma$ in a system $HSS[\emptyset]$ (homonymous synchronous systems), even if the membership is not known initially.*

Proof: From the definition of $H\Sigma$, it is enough to prove the following properties.

Validity. Since h_quora_p is a set, and the elements included in it are of the form $(mset, mset)$ (see Line 7 in Figure 7) there cannot be two pairs with the same label.

Monotonicity. The monotonicity of h_labels_p in Figure 7 holds because h_labels_p is initially empty, and each step, h_labels_p either grows or remains the same (see Line 8 in Figure 7). Similarly, the monotonicity of h_quora_p in Figure 7 follows from the fact that h_quora_p is initially empty, and any element $(mset, mset)$ included in it is never removed (see Line 7 in Figure 7).

Liveness. Let s be the synchronous step in which the last faulty process crashed. Then, in every step s' after s only correct processes will execute. Consider any process $p \in Correct$. In step s' will receive messages from all correct processes, and, hence, $mset_p = I(Correct)$. Then, process p includes $(I(Correct), I(Correct))$ in h_quora_p , and $I(Correct)$ in h_labels_p . Therefore, each correct process p is in $S(I(Correct))$. So, after step s , for each correct process p , the pair $(I(Correct), I(Correct))$ is in h_quora_p , and $I(Correct) = I(S(I(Correct))) \cap Correct$.

Safety. Consider two pairs $(x_1, x_1) \in h_quora_{p_1}^{\tau_1}$ and $(x_2, x_2) \in h_quora_{p_2}^{\tau_2}$, for any $p_1, p_2 \in \Pi$ and any $\tau_1, \tau_2 \in \mathbb{N}$.

Let M_1 be the set of processes from which p_1 received $(IDENT, -)$ messages in the synchronous step in which (x_1, x_1) was inserted for the first time in $h_quora_{p_1}$. Observe that $Correct \subseteq M_1$. Furthermore, any process $p \in S(x_1)$ must also be in M_1 (i.e., $S(x_1) \subseteq M_1$). Also, $x_1 = I(M_1)$, and, hence, $|x_1| = |M_1|$. Therefore, the only set $Q_1 \subseteq S(x_1)$ such that $I(Q_1) = x_1$ is $Q_1 = M_1$. We define M_2 similarly, and conclude that the only set $Q_2 \subseteq S(x_2)$ such that $I(Q_2) = x_2$ is $Q_2 = M_2$. Since $Q_1 \cap Q_2 \supseteq Correct \neq \emptyset$, the safety property holds. ■

```

1 operation propose( $v_p$ ):
2    $est1_p \leftarrow v_p; r_p \leftarrow 0$ ;
3   start Tasks T1 and T2;
4
5 Task T1
6   repeat forever
7      $r_p \leftarrow r_p + 1$ ;
8     // Leaders' Coordination Phase
9     broadcast ( $COORD, id(p), r_p, est1_p$ );
10    wait until ( $D.h.leader_p \neq id(p) \vee$ 
11      ( $D.h.multiplicity_p$  messages ( $COORD, id(p), r_p, -$ ) received));
12    if (some message ( $COORD, id(p), r_p, -$ ) received) then
13       $est1_p \leftarrow \min\{est_q : id(p) = id(q) \wedge$ 
14        ( $COORD, id(q), r_p, est_q$ ) received } end if;
15    // Phase 0
16    wait until ( $D.h.leader_p = id(p) \vee ((PH0, r_p, v)$  received);
17    if ( $(PH0, r_p, v)$  received) then  $est1_p \leftarrow v$  end if;
18    broadcast( $PH0, r_p, est1_p$ );
19    // Phase 1
20    broadcast( $PH1, r_p, est1_p$ );
21    wait until ( $PH1, r_p, -$ ) received from  $n - t$  processes;
22    if (the same estimate  $v$  received from  $> n/2$  processes) then
23       $est2_p \leftarrow v$ 
24    else
25       $est2_p \leftarrow \perp$ 
26    end if;
27    // Phase 2
28    broadcast( $PH2, r_p, est2_p$ );
29    wait until ( $PH2, r_p, -$ ) received from  $n - t$  processes;
30    let  $rec_p = \{est2 : \text{message } (PH2, r_p, est2) \text{ received}\}$ ;
31    if ( $(rec_p = \{v\}) \wedge (v \neq \perp)$ ) then
32      broadcast ( $DECIDE, v$ ); return( $v$ ) end if;
33    if ( $(rec_p = \{v, \perp\}) \wedge (v \neq \perp)$ ) then  $est1_p \leftarrow v$  end if;
34    if ( $rec_p = \{\perp\}$ ) then skip end if
35  end repeat.
36
37 Task T2
38 upon reception of ( $DECIDE, v$ ) do
39   broadcast ( $DECIDE, v$ ); return( $v$ ).

```

Figure 8: Consensus algorithm in $HAS[t < n/2, H\Omega]$ (code for process p). It uses detector $D \in H\Omega$.

5. Solving Consensus in Homonymous Systems

We present in this section two algorithms. One algorithm implements Consensus in $HAS[t < n/2, H\Omega]$, that is, in an homonymous asynchronous system with reliable links, using the failure detector $H\Omega$, and when a majority of processes are correct. The other algorithm implements Consensus in $HAS[H\Omega, H\Sigma]$, that is, in an homonymous asynchronous system with reliable links, using the failure detector $H\Omega$ and $H\Sigma$.

The main difficulty to solve in the two algorithms proposed is how to deal with the possibility of having multiple leaders, which is allowed by $H\Omega$. This is solved by adding to each round a preliminary phase in which homonymous leaders eventually “agree” in a common estimation of the value to propose. We call this additional phase Leaders’ Coordination Phase.

5.1. The Consensus Problem

In the Consensus problem, every process p proposes a value v_p and has to decide one value v_p^* , in such a way that the following properties are satisfied.

- Termination. Every correct process eventually decides.
- Validity. The value v_p^* decided by any process p is one of the proposed values.
- Agreement. All decided values are the same.

Each process p participates in the Consensus invoking the operation $\text{propose}(v_p)$. This operation returns to process p the value v_p^* it has decided.

5.2. Implementing Consensus in $HAS[t < n/2, H\Omega]$

Let us consider $HAS[t < n/2, H\Omega]$ where membership is unknown, but the number of processes is known (that is, n). Let us assume a majority of correct processes (i.e., $t < n/2$). We say that a process p is a leader, if it is correct and, after some finite time, $D.h_leader_q = id(p)$ permanently for each correct process q . By definition of $H\Omega$, there has to be at least one leader.

Brief explanation of the algorithm. The algorithm of Figure 8 is derived from the algorithm in Figure 4 of [4], proposed for anonymous systems. This algorithm has been adapted for homonymous systems. The algorithm of Figure 8 uses a failure detector of class $H\Omega$ (instead of $A\Omega$), and a new initial leaders’ coordination phase has been added. The purpose of this initial phase is to guarantee that, after a given round, all leaders propose the same value in each round.

The algorithm works in rounds, and it has four phases (Leaders’ Coordination Phase, Phase 0, Phase 1 and Phase 2). Every process p begins the Leaders’ Coordination phase of round r broadcasting a message $(COORD, id(p), r, est1_p)$. If process p considers itself a leader (querying the failure detector D of class $H\Omega$), it has to wait (Lines 10-11) until receiving messages $(COORD, id(p), r, est1)$ from its homonymous processes (whose number finds also querying the failure detector D of class $H\Omega$). After that, process p updates its estimate $est1_p$ with the minimal value proposed

among all its homonymous. Note that if p is a leader eventually all its homonymous will be leaders too. Hence, eventually all leaders will also choose the same minimal value in $est1$.

In Phase 0, if process p considers itself a leader (querying the failure detector D of class $H\Omega$) (Line 16), it broadcast a message $(PH0, r, est1_p)$ with its estimate in $est1_p$. Otherwise, process p has to wait until a message $(PH0, r, est1_l)$ is received from one of the leaders processes l , and update its variable $est1_p$ with the value received (Lines 16-17). Note that after the Leaders' Coordination Phase, eventually each leader l broadcasts messages $(PH0, r, est1_l)$ with the same value in $est1_l$.

In Phase 1, every process p broadcasts a message $(PH1, r, est1_p)$ with its estimate in $est1_p$. Then, it waits until the reception of messages $(PH1, r, est1)$ from a majority of processes. If the values $est1$ of all the received messages are equal (for example v), process p will adopt it, updating its variable $est2_p$ to v (Line 22). Otherwise, process p will update its variable $est2_p$ to \perp (Line 24). Note that when this Phase 1 finishes, the number of possible different values in the variables $est2_p$ of all processes are only two: v or \perp .

In Phase 2, every process p broadcasts a message $(PH2, r, est2_p)$ with its estimate in $est2_p$. Then, it waits until the reception of messages $(PH2, r, est2)$ from a majority of processes. If the value $est2$ of all the received messages is the same value v different from \perp , process p will decide v (Line 30). Otherwise, if some value $est2$ received is the value v different from \perp , process p will adopt it updating its variable $est1_p$ to v (Line 31) in order to propose it in the next round $r + 1$. With this, if a majority of processes (p not included) decides in this round r a value v , it will ensure that the process p will propose this same value v in the next round $r + 1$. Finally, if all the values $est2$ received from messages are the value \perp , nothing is performed.

Finally, Task 2 implements a reliable broadcast needed to propagate a decided value v from one process to the rest of processes of the system.

Correctness. The following lemmas are the key of the correctness of the algorithm. They show that, even having multiple leaders, these will eventually converge to propose the same value at each round.

Lemma 6. *No correct process blocks forever in any **wait** instruction in the algorithm of Figure 8.*

Proof: Let us consider by way of contradiction that the statement of the lemma is not correct. Hence, in some run of the algorithm of Figure 8 some correct process blocks forever in some **wait** instruction. Let us consider the smallest round r in which some correct process blocks permanently. Then, let us consider the **wait** instruction with the smallest number in which some correct process p blocks permanently in round r . Since there are four **wait** instructions in the algorithm, one in each phase, there are four cases to consider.

- Process p blocks forever in Lines 10-11 of the Leaders' Coordination Phase. Observe that p has to be a leader, because otherwise the first part of the wait condition is eventually satisfied and p would not block forever. Then, by definition of r , each leader q eventually reaches round r , and (even if it blocks in r) broadcasts $(COORD, id(q), r, -)$, where $id(q) = id(p)$, in Line 9. (Observe that all processes send $(COORD, -, -, -)$)

messages in Line 9, even if they do not consider themselves as leaders.) Eventually, $D.h_multiplicity_p$ holds permanently the number of leaders. Also eventually, all the $(COORD, id(q), r, -)$ messages sent by the leaders are delivered to p . Hence, the second part of the wait condition (Line 11) is satisfied. Thus, p is not blocked anymore, and, therefore, we reach a contradiction.

- Process p blocks forever in Line 16 of Phase 0. Observe that p cannot be a leader, because otherwise the first part of the wait condition is eventually satisfied and p would not block forever. For the same reason, no leader blocks forever in Line 16. Then, by definition of r and the fact that the first line in which any process blocks in round r is Line 16, each leader q eventually reaches Line 18 where it broadcasts $(PH0, r, -)$. Since there is at least one leader, at least one process sends such a message. Hence, p will eventually receive this message and the second part of the wait condition (Line 16) is satisfied. Thus, p is not blocked anymore, and, therefore, we reach a contradiction.
- Process p blocks forever in Line 21 of Phase 1. By definition of r and the fact that the first line in which any process blocks in round r is Line 21, each correct process q reaches Line 20 in round r . Then, each correct process q broadcasts $(PH1, r, -)$. Since there are at least $n - t$ correct processes by assumption, p will eventually receive $n - t$ messages $(PH1, r, -)$ and the condition of the wait instruction is satisfied. Thus, p is not blocked anymore, and, therefore, we reach a contradiction.
- Process p blocks forever in Line 29 of Phase 2. This case is similar to the previous one.

■

Lemma 7. *There is a round r such that at every round $r' > r$ all leaders broadcast the same value in Phase 0 of round r' , or there has been a decision in a round smaller than r' .*

Proof: Consider a time τ when all faulty processes have crashed and the failure detector D is stable (i.e., $\forall \tau' \geq \tau, \forall p \in Correct, D.h_leader_p^{\tau'} = \ell$, being $\ell \in I(Correct)$, and $D.h_multiplicity_p^{\tau'} = mult_{I(C)}(\ell)$). Let r be the largest round reached by any process at time τ . Then, we show that for any round $r' > r$, all leaders p have the same estimate $est1_p$ at the beginning of the Phase 0 of round r' (Line 16), or there has been a decision in a round smaller than r' . To prove this, let us assume that no decision is reached in a round smaller than r' . Then, since the leaders do not block forever in any round (Lemma 6), they execute Line 9 in round r' . Since the failure detector is stable, they also wait for the second part of the wait condition of Lines 10-11 (since the first part is not satisfied). When any leader p executes the Leaders' Coordination Phase of r' , it blocks in Lines 10-11 until it receives $D.h_multiplicity_p$ messages from the other leaders. By the stability of the $H\Omega$ failure detector, $D.h_multiplicity_p$ is the exact number of leaders. Also, from the definition of τ and r , no faulty process with identifier $D.h_leader_p$ is alive and all the messages they sent correspond to rounds smaller than r' . Hence, each leader p will wait to receive messages from all

the other leaders and will set $est1_p$ to the minimum from the same set of values (Line 14). ■

Using these lemmas and reusing some of the results in [4], we can conclude the correctness of the algorithm of Figure 8 in the following theorem.

Theorem 7. *The algorithm of Figure 8 solves consensus in $HAS[t < n/2, H\Omega]$.*

Proof: From the definition of Consensus, it is enough to prove the following properties.

Validity. The variable $est1$ is initialized with a value proposed by its process (Line 2). The value of $est1$ may be updated in Lines 14 or 17 with values of $est1$ broadcasted by other processes. The variable $est2$ is initialized and updated with $est1$ (Line 23) or \perp (Line 25). The value of $est1$ may be updated in Line 33 with values of $est2$ (different from \perp) broadcasted by other processes. The value decided in Line 32 is the value of $est2$ that was broadcasted by some process. As it is not possible to decide the value \perp (Line 32), then the value decided has to be one of the values proposed by the processes. Then, the validity property holds.

Agreement. Identical to the agreement property of Figure 4 of [4],

Termination. From Lemmas 6 and 7, after some round r , all leaders hold the same value v in $est1$ when they start executing Phase 0 of round r' (Line 16), and they broadcast this same value v (Line 18). Note that it is the same situation as having only one leader with value v stored in $est1$ when Phase 0 is reached. Hence, as Phase 0 starts in the same conditions as in the algorithm of Figure 4 of [4], the same proof can be used to prove the termination property. ■

5.3. Implementing Consensus in $HAS[H\Omega, H\Sigma]$

Figure 9 presents an algorithm that implements Consensus in $HAS[H\Omega, H\Sigma]$. Note that it is derived from the algorithm of Figure 4 of [3]⁷ where, like in the previous case, we have added a preliminary phase as a barrier such that homonymous leaders eventually “agree” in the same estimation value $est1$ to propose. Once this issue has been solved (as was proven for the previous algorithm), the use that this algorithm makes of the failure detector $H\Sigma$ is very similar to the use the algorithm of Figure 4 of [3] makes of the $A\Sigma$ failure detector.

Brief explanation of the algorithm. The algorithm of Figure 9 uses a failure detector of class $H\Omega$ (instead of $A\Omega$), a failure detector of class $H\Sigma$ (instead of $A\Sigma$), and a new initial leaders’ coordination Phase has been added. The purpose of this initial phase is to guarantee that, after a given round, all leaders propose the same value in each round. Note that this initial phase is analogous to the same phase used in the algorithm of Figure 8.

The algorithm works in rounds, and it has four phases (Leaders’ Coordination Phase, Phase 0, Phase 1 and Phase 2). Every process p begins the Leaders’ Coordination phase broadcasting a $(COORD, id(p), r, est1_p)$ message. If process p considers itself a leader (querying the failure detector $D1$ of class $H\Omega$), it has to wait until to receive

⁷Journal version in [6]

($COORD, id(p), r, est1$) messages sent by all its homonymous processes (also querying the failure detector $D1$ of class $H\Omega$) (Line 10). After that, process p updates its estimate $est1_p$ with the minimal value proposed among all its homonymous. Note that eventually all its homonymous will be leaders too. Hence, eventually all leaders will also choose the same minimal value in $est1$.

In Phase 0, if process p considers itself a leader (querying the failure detector $D1$ of class $H\Omega$) (Line 14), it broadcast a ($PH0, r, est1_p$) message with its estimate in $est1_p$. Otherwise, process p has to update its $est1_p$ waiting until a ($PH0, r, est1_l$) message is received from one of the leaders processes l (Lines 14-15). Note that after the Leaders' Coordination Phase, eventually each leader l broadcast ($PH0, -, est1_l$) messages with the same value in $est1_l$.

In Phase 1, every process p broadcasts a ($PH1, id(p), r, sr, current_labels_p, est1_p$) message, being r is the current round for process p , sr is the sub-round inside of Phase 1, $current_labels_p$ are the labels known by process p up to now. Process p waits for messages from some quorum of processes. Process p knows quora by reading the value of the variable $D2.h_quora_p$ of the failure detector $D2$ of class $H\Sigma$ (Line 22). The messages that form a quorum must satisfy the conditions of Lines 22-23. If when process p is building a quorum (i.e., when p is executing Lines 22-23) its failure detector variable $D2.h_labels_p$ in sr changes, then $current_labels_p$ is updated an a new sub-round starts, and process p broadcasts this new knowledge of labels. The same actions are taken if process p knows that another process has changed to an upper sub-round of sr (Lines 26-29). Phase 1 finishes when the messages that form a quorum finally satisfy the conditions of Lines 22-23. When this happens, if all the estimate values in the messages are equal (for example v), the variable $est2_p$ must be updated with this value v . Otherwise, $est2_p$ will be set to \perp (Lines 24-25).

In Phase 2, every process p broadcasts a ($PH2, id(p), r, sr, current_labels_p, est2_p$) message with its estimate in $est2_p$. Then, process p waits for messages from some quorum of processes to take a decision. This Phase 2 is very similar to Phase 1 (it runs in sub-rounds, which are increased when process p knows that its variable $D2.h_labels_p$ changes). Similarly to Phase 1, when the messages that form a quorum finally satisfy the conditions of Lines 37-38, if all the estimate values in the messages are equal (for example $v \neq \perp$), process p will decide v (Line 40). Otherwise, if some of the estimate $est2$ received from messages are the value v different from \perp , process p will take it updating its variable $est1_p$ to v (Line 41) in order to propose it in the next round $r + 1$. Finally, if all estimate $est2$ received from messages are the value \perp , nothing is performed.

Task 2 implements a reliable broadcast needed to propagate a decided value v from one process to the rest of processes of the system.

Correctness. In order to prove the correctness of the algorithm we start by proving the following lemmas.

Lemma 8. *No correct process blocks forever in the repeat loops of Phases 1 and 2.*


```

1 operation propose( $v_p$ ):
2  $est1_p \leftarrow v_p$ ;  $r_p \leftarrow 0$ ;
3 start Tasks T1 and T2;
4
5 Task T1
6 repeat forever
7    $r_p \leftarrow r_p + 1$ ;
8   // Leaders' Coordination Phase
9   broadcast ( $COORD, id(p), r_p, est1_p$ );
10  wait until ( $D1.h.leader_p \neq id(p)$ )  $\vee$ 
11    ( $D1.h.multiplicity_p$  messages ( $COORD, id(p), r_p, -$ ) received);
12  if (some message ( $COORD, id(p), r_p, -$ ) received) then
13     $est1_p \leftarrow \min\{est_q : id(p) = id(q) \wedge$ 
14      ( $COORD, id(q), r_p, est_q$ ) received } end if;
15  // Phase 0
16  wait until ( $D1.h.leader_p = id(p) \vee ((PH0, r_p, v)$  received);
17  if ( $(PH0, r_p, v)$  received) then  $est1_p \leftarrow v$  end if;
18  broadcast ( $PH0, r_p, est1_p$ );
19  // Phase 1
20   $sr_p \leftarrow 1$ ;  $current.labels_p \leftarrow D2.h.labels_p$ ;
21  broadcast ( $PH1, id(p), r_p, sr_p, current.labels_p, est1_p$ );
22  repeat
23    if ( $(PH2, -, r_p, -, -, est2)$  received) then
24       $est2_p \leftarrow est2$ ; exit inner repeat loop end if;
25    if ( $(\exists(x, mset) \in D2.h.quora_p) \wedge (\exists sr \in \mathbb{N}) \wedge$ 
26      ( $\exists$  set  $M$  of messages ( $PH1, -, r_p, sr, -, -$ )), such that,
27      ( $\forall(PH1, -, -, -, cl, -) \in M, x \in cl) \wedge$ 
28      ( $mset = \{i : (PH1, i, -, -, -, -) \in M\}$ )) then
29      if (all msgs in  $M$  contain the same estimate  $v$ ) then
30         $est2_p \leftarrow v$  else  $est2_p \leftarrow \perp$  end if;
31      exit inner repeat loop;
32    else if ( $current.labels_p \neq D2.h.labels_p$ )  $\vee$ 
33      ( $(PH1, -, r_p, sr, -, -)$  recvd with  $sr > sr_p$ ) then
34       $sr_p \leftarrow sr_p + 1$ ;  $current.labels_p \leftarrow D2.h.labels_p$ ;
35      broadcast ( $PH1, id(p), r_p, sr_p, current.labels_p, est1_p$ );
36    end if
37  end if
38  end repeat;
39  // Phase 2
40   $sr_p \leftarrow 1$ ;  $current.labels_p \leftarrow D2.h.labels_p$ ;
41  broadcast ( $PH2, id(p), r_p, sr_p, current.labels_p, est2_p$ );
42  repeat
43    if ( $(COORD, -, r_p + 1, -)$  received) then
44      exit inner repeat loop end if;
45    if ( $(\exists(x, mset) \in D2.h.quora_p) \wedge (\exists sr \in \mathbb{N}) \wedge$ 
46      ( $\exists$  set  $M$  of messages ( $PH2, -, r_p, sr, -, -$ )), such that,
47      ( $\forall(PH2, -, -, -, cl, -) \in M, x \in cl) \wedge$ 
48      ( $mset = \{i : (PH2, i, -, -, -, -) \in M\}$ )) then
49      let  $rec_p$  = the set of estimates contained in  $M$ ;
50      if ( $(rec_p = \{v\}) \wedge (v \neq \perp)$ ) then
51        broadcast ( $DECIDE, v$ ); return( $v$ ) end if;
52      if ( $(rec_p = \{v, \perp\}) \wedge (v \neq \perp)$ ) then  $est1_p \leftarrow v$  end if;
53      if ( $rec_p = \{\perp\}$ ) then skip end if;
54      exit inner repeat loop
55    else if ( $(current.labels_p \neq D2.h.labels_p) \vee$ 
56      ( $(PH2, -, r_p, sr, -, -)$  received with  $sr > sr_p$ )) then
57       $sr_p \leftarrow sr_p + 1$ ;  $current.labels_p \leftarrow D2.h.labels_p$ ;
58      broadcast ( $PH2, id(p), r_p, sr_p, current.labels_p, est2_p$ );
59    end if
60  end if
61  end repeat
62  end repeat.
63
64 Task T2
65 upon reception of ( $DECIDE, v$ ) do
66   broadcast ( $DECIDE, v$ ); return( $v$ ).

```

Figure 9: Consensus algorithm in $HAS[H\Omega, H\Sigma]$ (code for process p). It uses detectors $D1 \in H\Omega$ and $D2 \in H\Sigma$.

Proof: Note that if a correct process decides (Line 51), then the claims follows. Consider the repeat loop of Phase 1 (Lines 22-38). Let us assume that some correct process is blocked forever in this loop. Then, let us consider the first round r in which a correct process blocks forever in r . Hence, all correct processes must block forever in the same loop in round r . Otherwise some process broadcasts a message $(PH2, -, r, -, -, -)$, and from Line 24 no correct process would block forever in this loop of round r . Let us consider a correct process p , and the pair (x, m) that guarantees the liveness property for p . Then, there is a time in which $(x, m) \in D2.h_quora_p$ and every correct process q in $S(x) \cap Correct$ has $x \in D2.h_labels_q$. Note that, from Lines 32-36, every change in the variable $D2.h_labels$ of a process creates a new sub-round, and that all processes broadcast their current value of $D2.h_labels$ in each new sub-round. Therefore, eventually, p will receive messages $(PH1, -, r, sr, cl, -)$ from all these processes such that $x \in cl$. Hence, the condition of Lines 25-28 is satisfied, and p will exit the loop of Phase 1. The argument for the repeat loop of Phase 2 is verbatim. ■

Lemma 9. *No two processes decide different values in the same round.*

Proof: Let us assume that processes p_1 and p_2 decide values v_1 and v_2 in sub-rounds sr_1 and sr_2 , respectively, of the same round r (in Line 51). Let (x_1, m_1) and M_1 be the pair in $D2.h_quora_{p_1}$ and the set of messages that satisfy the condition of Lines 45-48 for p_1 . Since for each message $(PH2, -, r, sr_1, cl, -) \in M_1$, it holds that $x_1 \in cl$, if Q_1 is the set of senders of the messages in M_1 , we have that $Q_1 \subseteq S(x_1)$. Additionally, $m_1 = \{i : (PH2, i, -, -, -, -) \in M_1\} = I(Q_1)$. We can define (x_2, m_2) and M_2 analogously for p_2 . Then, from the Safety Property of $H\Sigma$, $Q_1 \cap Q_2 \neq \emptyset$. Let $p_l \in Q_1 \cap Q_2$. Then, process p_l must have broadcast messages $(PH2, id(p_l), r, sr_1, -, v_1)$ and $(PH2, id(p_l), r, sr_2, -, v_2)$ (Lines 41 and 58). Since the estimate $est2_{p_l}$ of p_l does not change between sub-rounds (inner repeat loop, Lines 42-61), it must hold that $v_1 = v_2$. From the condition of Line 51, $rec_{p_1} = \{v_1\}$ in sub-round sr_1 and $rec_{p_2} = \{v_2\}$ in sub-round sr_2 , and both processes decide the same value. Hence, no two processes decide different values in the same round. ■

Using the above lemmas and reusing some of the results from [3] the correctness of the algorithm of Figure 9 can be shown.

Theorem 8. *The algorithm of Figure 9 solves consensus in $HAS[H\Omega, H\Sigma]$.*

Proof: The proof of this theorem is similar to the proof of Theorem 5 of [3], with the following changes. Observe that the Leaders' Coordination Phase and Phase 0 of the algorithms in Figures 8 and 9 are the same. Hence, Lemmas 6 and 7 also apply to the algorithm of Figure 9. Then, the termination property can be proven in a similar way as in [3] (Lemmas 1 and 2), but using those two Lemmas 6 and 7 together with Lemma 8. The proof of the agreement property is also similar to Lemma 3 of [3] but using Lemma 9. ■

Observe that the algorithm of Figure 9 can be easily transformed into an algorithm that solves consensus in $AAS[A\Omega, H\Sigma]$ (an anonymous system with detectors $A\Omega$ and $H\Sigma$). For that, given a failure detector $D3 \in A\Omega$, it is enough to remove the Leaders' Coordination Phase, and in Phase 0 to replace $(D1.h_leader_p = id(p))$ by $(D3.a_leader_p)$. The resulting Phase 0 is the same as Phase 1 in the algorithm of Figure 3 of [6], and has the same properties.

6. Conclusion

This paper studied the Consensus problem in a new distributed environment with homonymous processes. New failure detectors have been defined for this homonymous model (called $H\Omega$, $\diamond H\bar{P}$ and $H\Sigma$). We have studied the relations among the failure detector classes Σ and its versions for homonymous systems (denoted by $H\Sigma$), and for anonymous systems (denoted by $A\Sigma$). It has also been shown that $H\Omega$, $\diamond H\bar{P}$ and $H\Sigma$ can be implemented in homonymous systems with different synchrony requirements (and in all cases without initial knowledge of the membership). Interestingly, our class $H\Omega$ can be implemented with partial synchrony, while the analogous class $A\Omega$ defined for anonymous systems can not be implemented (even in synchronous systems). Finally, we have shown that $H\Omega$ (and a majority of correct processes), and $\langle H\Omega, H\Sigma \rangle$ failure detectors can be used to implement Consensus in homonymous asynchronous systems (even without initial knowledge of the membership).

References

- [1] D. Angluin. Local and global properties in networks of processors (extended abstract). In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing (STOC)*, 28-30 April, Los Angeles, California, pages 82–93. ACM, 1980.
- [2] H. Attiya, M. Snir, and M. Warmuth. Computing on an anonymous ring. *J. ACM*, 35(4):845–875, 1988.
- [3] F. Bonnet and M. Raynal. Anonymous asynchronous systems: The case of failure detectors. Technical Report PI 1945, IRISA, Rennes, France, January 2010.
- [4] F. Bonnet and M. Raynal. Consensus in anonymous distributed systems: Is there a weakest failure detector? In *Proceedings of 24th IEEE International Conference on Advanced Information Networking and Applications (AINA)*, 2010, Perth, Australia, 20-13 April 2010, pages 206–213. IEEE Computer Society, 2010.
- [5] F. Bonnet and M. Raynal. The price of anonymity: Optimal consensus despite asynchrony, crash, and anonymity. *ACM Transactions on Autonomous and Adaptive Systems, TAAS*, 6(4):23, 2011.
- [6] François Bonnet and Michel Raynal. Anonymous asynchronous systems: the case of failure detectors. *Distributed Computing*, 26(3):141–158, 2013.

- [7] Z. Bouzid, P. Sutra, and C. Travers. Anonymous agreement: The janus algorithm. In Antonio Fernández Anta, Giuseppe Lipari, and Matthieu Roy, editors, *OPODIS*, volume 7109 of *Lecture Notes in Computer Science*, pages 175–190. Springer, 2011.
- [8] Z. Bouzid and C. Travers. Brief announcement: Anonymity, failures, detectors and consensus. In M. Aguilera, editor, *DISC*, volume 7611 of *Lecture Notes in Computer Science*, pages 427–428. Springer, 2012.
- [9] T. Chandra, V. Hadzilacos, and S. Toueg. The weakest failure detector for solving consensus. *J. ACM*, 43(4):685–722, 1996.
- [10] T. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. *J. ACM*, 43(2):225–267, 1996.
- [11] C. Delporte-Gallet, H. Fauconnier, and R. Guerraoui. A realistic look at failure detectors. In *Proceedings of International Conference on Dependable Systems and Networks (DSN 2002), 23-26 June 2002, Bethesda, MD, USA*, pages 345–353. IEEE Computer Society, 2002.
- [12] C. Delporte-Gallet, H. Fauconnier, and R. Guerraoui. Tight failure detection bounds on atomic object implementations. *J. ACM*, 57(4), 2010.
- [13] C. Delporte-Gallet, H. Fauconnier, and A. Tielmann. Fault-tolerant consensus in unknown and anonymous networks. In *Proceedings of 29th IEEE International Conference on Distributed Computing Systems (ICDCS 2009), 22-26 June 2009, Montreal, Québec, Canada*, pages 368–375. IEEE Computer Society, 2009.
- [14] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, Anne-Marie Kermarrec, Eric Ruppert, et al. Byzantine agreement with homonyms. *Distributed Computing*, 26(3), 2013.
- [15] M. Fischer, N. Lynch, and M. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.
- [16] F. Greve and S. Tixeuil. Knowledge connectivity vs. synchrony requirements for fault-tolerant agreement in unknown networks. In *Proceedings of 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2007), 25-28 June 2007, Edinburgh, UK, Proceedings*, pages 82–91. IEEE Computer Society, 2007.
- [17] E. Jiménez, S. Arévalo, and A. Fernández. Implementing unreliable failure detectors with unknown membership. *Inf. Process. Lett.*, 100(2):60–63, 2006.
- [18] N. Lynch. *Distributed Algorithms*. Morgan Kaufmann Pub., San Francisco (CA), 1996.
- [19] M. Raynal. Failure detectors for asynchronous distributed systems: an introduction. In *Wiley Encyclopedia of Computer Science and Engineering*, volume 2, pages 1181–1191. 2009.

- [20] M. Raynal. *Communication and Agreement Abstractions for Fault-Tolerant Asynchronous Distributed Systems*. Morgan & Claypool Publishers, 2010.
- [21] M. Yamashita and T. Kameda. Computing on anonymous networks: Part i-characterizing the solvable cases. *IEEE Trans. Parallel Distrib. Syst.*, 7(1):69–89, 1996.
- [22] P. Zielinski. Anti-omega: the weakest failure detector for set agreement. In R. Bazzi and B. Patt-Shamir, editors, *PODC*, pages 55–64. ACM, 2008.