



HAL
open science

Automatic Transformations from Crash-Stop to Permanent Omission

Carole Delporte-Gallet, Hugues Fauconnier, Felix Freiling, Lucia Draque
Penso, Andreas Tielmann

► **To cite this version:**

Carole Delporte-Gallet, Hugues Fauconnier, Felix Freiling, Lucia Draque Penso, Andreas Tielmann.
Automatic Transformations from Crash-Stop to Permanent Omission. 2007. hal-00160626v2

HAL Id: hal-00160626

<https://hal.science/hal-00160626v2>

Preprint submitted on 10 Oct 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

From Crash-Stop to Permanent Omission: Automatic Transformation and Weakest Failure Detectors

Carole Delporte-Gallet¹, Hugues Fauconnier¹, Felix C. Freiling²,
Lucia Draque Penso², Andreas Tielmann¹

¹ University of Paris 7 - Denis Diderot

Laboratoire d'Informatique Algorithmique: Fondements et Applications (LIAFA)

² University of Mannheim

Laboratory for Dependable Distributed Systems

September 2007

Abstract—This paper studies the impact of omission failures on asynchronous distributed systems with crash-stop failures. We provide two different transformations for algorithms, failure detectors, and problem specifications, one of which is weakest failure detector preserving. We prove that our transformation of failure detector Ω [1] is the weakest failure detector for consensus in environments with crash-stop and permanent omission failures and a majority of correct processes. Our results help to use the power of the well-understood crash-stop model to automatically derive solutions for the general omission model, which has recently raised interest for being noticeably applicable for security problems in distributed environments equipped with security modules such as smartcards [2–4].

Index Terms—Fault-Tolerance, Weakest Failure Detectors, Transformations, Asynchronous Systems, Crash-Stop, Permanent Omissions.

I. INTRODUCTION

MESSAGE omission failures, which have been introduced by Hadzilacos [5] and been refined by Perry and Toueg [6], put the blame of a message loss to a specific process instead of an unreliable message channel. Beyond the theoretical interest, omission models are also interesting for practical problems like they arise from the security area: Assume that some kind of trusted smartcards are disposed on untrusted processors. If these smartcards execute trusted algorithms and are able to sign messages, then it is relatively easy to restrict the power of a malicious adversary to only be able to drop messages of the trusted smartcards or to stop the smartcards themselves. Following this approach, omission models have lead to the development of reductions from security problems in the Byzantine failure model [7] such as fair-exchange [3, 4], and secure multiparty computation [2] to well-known distributed problems in the general omission model, such as consensus [8], where both process crashes and message omissions may take place. Apart from that,

omission failures can model overflows of local message buffers in typical communication environments.

The message omission and crash failures are considered here in asynchronous systems. Due to classical impossibility results concerning problems as consensus [9] in asynchronous systems, following the failure detector approach [10], we augment the system with oracles that give information about failures.

The extension of failure detectors to more severe failure models than crash failures is unclear [11], because in these models failures may depend on the scheduling and on the algorithm. As it is easy to transform the general omission model into a model with only permanent omissions using standard techniques like the piggybacking of messages, we consider only permanent omissions and crashes. This means that if an omission failure occurs, then it occurs permanently. In this model, precise and simple definitions for failure detectors can easily be deduced from the ones in the crash-stop model.

To provide the permanent omission model with the benefits of a well-understood system model like the crash-stop model, we give automatic transformations for problem specifications, failure detectors, and algorithms such that algorithms designed to tolerate only crash-stop failures can be executed in permanent omission environments and use transformed failure detectors to solve transformed problems. Specifically, we give two transformations. At first, one that works in every environment, but that transforms uniform problems into problems with only limited uniformity, and at second one that works only with a majority of correct processes, but transforms uniform crash-stop problems into their uniform permanent omission counterpart. An interesting point is the fact that the transformation of the specification gives for most of the classical problems the standard specification in the message omission and crash failure model. For example, from an algorithmic solution A of the consensus problem with a failure detector \mathcal{D} in the crash-stop model, we automatically get $A' = \text{trans}(A)$, an algorithmic solution of the consensus

problem using $\mathcal{D}' = \text{trans}(\mathcal{D})$ in the message omission and crash failure model.

Moreover, our first transformation preserves also the “weaker than” relation [1] between failure detectors. This means that if a failure detector is a weakest failure detector for a certain (crash-stop) problem, then its transformation is a weakest failure detector for the transformed problem. We can use this to show that our transformation of failure detector Ω [1] is the weakest failure detector for (uniform) consensus in an environment with permanent omission failures and a majority of correct processes.

The problem of automatically increasing the fault-tolerance of algorithms in environments with crash-stop failures has been extensively studied before [12–15]. The results of Neiger and Toueg [13], Delporte-Gallet et al. [14], and Bazzi and Neiger [15] assume in contrast to ours synchronous systems and no failure detectors. Neiger and Toueg [13] propose several transformations from crash-stop to send omission, to general omission, and to Byzantine faults. Delporte-Gallet et al. [14] transform round-based algorithms with broadcast primitives into crash-stop-, general omission-, and Byzantine-tolerant algorithms. Asynchronous systems are considered by Basu, Charron-Bost, and Toueg [12] but in the context of link failures instead of omission failures and also without failure detectors. The types of link failures that are considered by Basu, Charron-Bost, and Toueg [12] are eventually reliable and fair-lossy links. Eventually reliable links can lose a finite (but unbounded) number of messages and fair-lossy links satisfy that if infinitely many messages are sent over it, then infinitely many messages do not get lost. To show our results, we extend the system model of Basu, Charron-Bost, and Toueg [12] such that we can model omission failures, failure patterns, and failure detectors. Another definition for a system model with crash-recovery failures, omission failures, and failure detectors is given by Dolev et al. [16]. In this model, the existence of a fully connected component of processes that is completely detached from all other processes is assumed and only the processes in this component are declared to be correct.

The omission failure detector defined by Delporte-Gallet et al. [17] that can be implemented in partially synchronous models using some weak timing assumptions, is in comparison with our transformed Ω strictly stronger. However, with a correct majority, both failure detectors can easily be transformed into each other.

To the best of our knowledge, this is the first paper that investigates an automatic transformation to increase the fault tolerance of distributed algorithms in asynchronous systems augmented with failure detectors.

We organize this paper as follows. In Section II, we define our formal system model, in Section III, we define our general problem and algorithm transformations, in Section IV we state and prove our theorems, and finally, in Section V, we summarize and discuss our results.

II. MODEL

The asynchronous distributed system is assumed to consist of n distinct fully-connected processes $\Pi = \{p_1, \dots, p_n\}$. The

asynchrony of the system means, that there are no bounds on the relative process speeds and message transmission delays. To allow an easier reasoning, a discrete global clock \mathcal{T} is added to the system. The system model used here is derived from that of Basu, Charron-Bost, and Toueg [12]. It has been adapted to model also failure detectors and permanent omission failures.

1) *Algorithms*: An algorithm A is defined as a vector of local algorithm modules (or simply modules) $A(\Pi) = \langle A(p_1), \dots, A(p_n) \rangle$. Each local algorithm module $A(p_i)$ is associated with a process $p_i \in \Pi$ and defined as a deterministic infinite state automaton. The local algorithm modules can exchange messages via send and receive primitives. We assume all messages to be unique.

2) *Failures and Failure Patterns*: A failure pattern \mathcal{F} is a function that maps each value t from \mathcal{T} to an output value that specifies which failures have occurred up to time t during an execution of a distributed system. Such a failure pattern is totally independent of any algorithm. A crash-failure pattern

$$C : \mathcal{T} \rightarrow 2^\Pi$$

denotes the set of processes that have crashed up to time t ($\forall t : C(t) \subseteq C(t+1)$).

Additionally to the crash of a process, it can fail by not sending or not receiving a message. We say that it *omits* a message. The message omissions do not occur because of link failures, they model overflows of local message buffers or the behavior of a malicious adversary with control over the message flow of certain processes. It is important that for every omission, there is a process responsible for it. As we already mentioned, we consider only permanent omissions and leave the treatment of transient omissions over to the underlying asynchronous communication layer. Intuitively, a process has a permanent send omission if it always fails by not sending messages to a certain other process after a certain point in time. Analogously, a process has a permanent receive omission if it always fails by not receiving messages from a certain other process after a certain point in time. The permanent omissions are modeled via a send- and a receive-omission failure pattern:

$$O_S : \mathcal{T} \rightarrow 2^{\Pi \times \Pi} \quad \text{and} \quad O_R : \mathcal{T} \rightarrow 2^{\Pi \times \Pi}$$

If $(p_s, p_d) \in O_S(t)$, then process p_s has a permanent send-omission to process p_d after time t . If $(p_s, p_d) \in O_R(t)$, then process p_d has a permanent receive-omission to process p_s after time t . All the failure patterns defined so far can be put together to a single failure pattern $\mathcal{F} = (C, O_S, O_R)$.

With such a failure pattern, we define a process to be *correct*, if it experiences no failure at all. We assume that at least one process is correct. A process p is crash-correct ($p \in \text{cr-correct}(\mathcal{F})$) in \mathcal{F} , if it does not crash.

A process p_d is *directly-reachable* from another process p_s in \mathcal{F} , if for all $t \in \mathcal{T}$, $(p_s, p_d) \notin O_S(t)$ and $(p_s, p_d) \notin O_R(t)$. A process p_d is called *reachable* from a process p_s , if p_d is directly-reachable from p_s , or if there exists a process q , such that p_d is reachable from q and q is reachable from p_s (transitive closure). If a process is reachable from some correct processes, then it is *in-connected*. Analogously, a process is *out-connected*, if some correct processes are reachable from

it. If a process p is in-connected and out-connected in a failure pattern \mathcal{F} , then we say that p is *connected* in \mathcal{F} ($p \in \text{connected}(\mathcal{F})$). This means that between connected processes there is always reliable communication possible. With a simple relaying algorithm, every message can eventually be delivered. Note that it is nevertheless still possible that connected processes receive messages from disconnected processes or disconnected processes receive messages from connected ones. The difference between connected and disconnected processes is that the former are able to send and to receive messages to/from correct processes and therefore are able to communicate in both directions. It is easy to see that $\text{crash-correct}(\mathcal{F}) \supseteq \text{connected}(\mathcal{F}) \supseteq \text{correct}(\mathcal{F})$.

We say that a failure pattern \mathcal{F}' is an *omission equivalent extension* of another failure pattern \mathcal{F} ($\mathcal{F} \leq_{om} \mathcal{F}'$), if the set of crash-correct processes in \mathcal{F} is at all times equal to the set of connected processes in \mathcal{F}' and there are no omission failures in \mathcal{F} .

We define an *environment* \mathcal{E} to be a set of possible failure patterns. $\mathcal{E}_{c.s.}^f$ denotes the set of all failure patterns where only crash-stop faults occur and at most f processes crash. $\mathcal{E}_{p.o.}^f$ denotes the set of all failure patterns where crash-stop and permanent omission faults may occur and at most f processes are not connected (clearly, $\mathcal{E}_{c.s.}^f \subseteq \mathcal{E}_{p.o.}^f$).

3) *Failure Detectors*: A failure detector provides (possibly incorrect) information about a failure pattern [10]. Associated with each failure detector is a (possibly infinite) range \mathcal{R} of values output by that failure detector. A failure detector history FDH with range \mathcal{R} is a function from $\Pi \times \mathcal{T}$ to \mathcal{R} . $FDH(p, t)$ is the value of the failure detector module of process p at time t . A failure detector \mathcal{D} is a function that maps a failure pattern \mathcal{F} to a *set* of failure detector histories with range \mathcal{R} . $\mathcal{D}(\mathcal{F})$ denotes the set of possible failure detector histories permitted by \mathcal{D} for the failure pattern \mathcal{F} . Note that a failure detector \mathcal{D} is specified as a function of the failure pattern \mathcal{F} of an execution. However, an implementation of \mathcal{D} may use other aspects of the execution such as when messages are arrived and executions with the same failure pattern \mathcal{F} may still have different failure detector histories. It is for this reason that we allow $\mathcal{D}(\mathcal{F})$ to be a set of failure detector histories from which the actual failure detector history for a particular execution is selected non-deterministically.

Take failure detector Ω [1] as an example. The output of the failure detector module of Ω at a process p_i is a *single* process, p_j , that p_i currently considers to be *crash-correct*. In this case, the range of output values is $\mathcal{R}_\Omega = \Pi$. For each failure pattern \mathcal{F} , $\Omega(\mathcal{F})$ is the set of all failure detector histories FDH_Ω with range \mathcal{R}_Ω that satisfy the following property: There is a time after which all the crash-correct processes always trust the same crash-correct process:

$$\begin{aligned} \exists t \in \mathcal{T}, \exists p_j \in \text{cr-correct}(\mathcal{F}), \\ \forall p_i \in \text{cr-correct}(\mathcal{F}), \forall t' \geq t : FDH_\Omega(p_i, t') = p_j \end{aligned}$$

The output of failure detector module Ω at a process p_i may change with time, i.e. p_i may trust different processes at different times. Furthermore, at any given time t , processes p_i and p_j may trust different processes.

A local algorithm module $A(p_i)$ can access the current output value of its local failure detector module using the action *queryFD*.

4) *Histories*: A local history of a local algorithm module $A(p_i)$, denoted $H[i]$, is a finite or an infinite sequence of alternating states and events of type *send*, *receive*, *queryFD*, or *internal*. We assume that there is a function *time* that assigns every event to a certain point in time and define $H[i]/_t$ to be the maximal prefix of $H[i]$ where all events have occurred before time t . A *history* H of $A(\Pi)$ is a vector of local histories $\langle H[1], H[2], \dots, H[n] \rangle$.

5) *Reliable Links*: A reliable link does not create, duplicate, or lose messages. Specifically, if there is no permanent omission between two processes and the recipient executes infinitely many receive actions, then it will eventually receive every message. We specify, that our underlying communication channels ensure reliable links.

6) *Problem Specifications*: Let Π be a set of processes and A be an algorithm. We define $\mathcal{H}(A(\Pi), \mathcal{E})$ to be the set of all tuples (H, \mathcal{F}) such that H is a history of $A(\Pi)$, $\mathcal{F} \in \mathcal{E}$, and H and \mathcal{F} are compatible, that is crashed processes do not take any steps after the time of their crash, there are no receive-events after a permanent omission, etc. A *system* $\mathcal{S}(A(\Pi), \mathcal{E})$ of $A(\Pi)$ is a subset of $\mathcal{H}(A(\Pi), \mathcal{E})$. A *problem specification* Σ is a set of tuples of histories and failure patterns, because (permanent) omission failures are not necessarily reflected in a history (e.g., if a process sends no messages). A system \mathcal{S} satisfies a problem specification Σ , if $\mathcal{S} \subseteq \Sigma$. We say that an algorithm A satisfies a problem specification Σ in environment \mathcal{E} , if $\mathcal{H}(A(\Pi), \mathcal{E}) \subseteq \Sigma$.

Take consensus as an example (see Table I): It is specified by making statements about some variables *propose* and *decide* in the states of a history (e.g. the value of *decide* has eventually to be equal at all (crash-)correct processes). This can be expressed as the set of all tuples (H, \mathcal{F}) where there exists a time t and a value v , such that for all $p_i \in \text{cr-correct}(\mathcal{F})$, there exists an event e in $H[i]$ with $\text{time}(e) \leq t$ and for all states s after event e , the value of the variable *decide* in s is v .

III. FROM CRASH-STOP TO PERMANENT OMISSION

We will give here two transformations: one general transformation for all environments, where we provide only restricted guarantees for disconnected processes, and one for environments where less than half of the processes may not be connected, where we are able to provide for all processes the same guarantees as for the crash-stop case.

To improve the fault-tolerance of algorithms, we simulate a single state of the original algorithm with several states of the simulation algorithm. For these additional states, we *augment* the original states with additional variables. Since an event of the simulation algorithm may lead to a state where only the augmentation variables change, the sequence of the original variables may *stutter*. We call a local history $H'[i]$ a *stuttered* and *augmented extension* of a history $H[i]$ ($H[i] \leq_{sa} H'[i]$), if $H[i]$ and $H'[i]$ differ only in the value of the augmentation variables and some additional states caused by differences in

these variables (in particular, $H[i] \leq_{sa} H'[i]$ for all $H[i]$). If $H[i] \leq_{sa} H'[i]$ for all $p_i \in \Pi$, we write $H \leq_{sa} H'$. We say that a problem specification Σ is closed under stuttering and augmentation, if $(H, \mathcal{F}) \in \Sigma$ and $H \leq_{sa} H'$ implies that (H', \mathcal{F}) is also in Σ . Most problems satisfy this natural closure property (e.g. consensus).

A. The General Transformation

1) *Transformation of Problem Specifications:* To transform a problem specification, we first show a transformation of a tuple of a trace and a failure pattern. Based on this transformation, we transform a whole problem specification. The intuition behind this transformation is that we demand only something from processes as long as they are connected. After their disconnection, processes may behave arbitrary. More formally, let $t_{c.s.}(i)$ be the time at which process p_i crashes in \mathcal{F} ($t_{c.s.}(i) = \infty$, if p_i never crashes). Analogously, let $t'_{p.o.}(i)$ be the time at which process p_i becomes disconnected in \mathcal{F}' ($t'_{p.o.}(i) = \infty$, if p_i never becomes disconnected). Then:

$$(H', \mathcal{F}') \in \text{trans}((H, \mathcal{F})) \\ \Leftrightarrow \forall p_i \in \Pi : H[i]/t_{c.s.}(i) \leq_{sa} H'[i]/t'_{p.o.}(i)$$

and for a whole problem specification:

$$\text{trans}(\Sigma) := \{(H', \mathcal{F}') \mid \\ (H', \mathcal{F}') \in \text{trans}((H, \mathcal{F})) \wedge (H, \mathcal{F}) \in \Sigma\}$$

A transformation of non-uniform consensus, where properties of certain propose- and decision-variables of (crash-)correct processes are specified would lead to a specification where the same properties are ensured for the states of connected processes, because only histories with the same states (disregarding the augmentation variables) are allowed in the transformation at this processes (see Table I). We also take the states of processes *before* they become disconnected into account, because they (e.g. their initial states for the propose variables) may also have an influence on the fulfillment of a problem specification, although they are after their disconnection not allowed to have this influence anymore. Since we impose no restriction on the behavior of processes after their disconnection, the transformed problem specification allows them to decide a value that was never proposed (although our transformation algorithms guarantee that this will not happen).

A transformation of uniform consensus leads to a problem specification where the uniform agreement is only demanded for processes before their time of disconnection. This means that it is allowed that after a partitioning of the network, the processes in the different network partitions come to different decision values. Another transformation, in which uniform consensus remains truly uniform is given in Section III-B.

2) *Transformation of Failure Detector Specifications:* We allow all failure detector histories for a failure pattern \mathcal{F} in $\text{trans}(\mathcal{D})$ that are allowed in the crash-stop version \mathcal{F}' of \mathcal{F} in \mathcal{D} :

$$\text{trans}(\mathcal{D})(\mathcal{F}) := \bigcup_{\mathcal{F}'} \{\mathcal{D}(\mathcal{F}') \mid \mathcal{F}' \leq_{om} \mathcal{F}\}$$

Consider failure detector Ω [1]. Ω outputs only failure detector histories that eventually provide the same crash-correct leader at all crash-correct processes. Then, $\text{trans}(\Omega)$ outputs these failure detector histories if and only if they provide a *connected* common leader at all *connected* processes.




Fig. 1. Additional Communication Layers

3) *Transformation of Algorithms:* In our algorithm transformation, we add new communication layers such that some of the omission failures in the system become transparent to the algorithm (see Figure 1). We transform a given algorithm A into another algorithm $A' = \text{trans}(A)$ in two steps:

- In the first step, we remove the send and receive actions from A and simulate them with a *three-way-handshake (3wh) algorithm*. The algorithm is described in Figure 2. The idea of the 3wh-algorithm is to substitute every send-action with an exchange of three messages. This means that to send a message to a certain process, it is necessary for a process to be able to send *and* to receive messages from it. Moreover, while the communication between connected processes is still possible, processes that are only in-connected or only out-connected (and not both) become totally disconnected. Hence, we eliminate influences of disconnected processes not existing in the crash-stop case.
- Then, in the second step, we remove the send and receive actions from the three way handshake algorithm and simulate them with a *relaying algorithm*. The idea of the relay algorithm is to relay every message to all other processes, such that they relay it again and all connected processes can communicate with each other, despite the fact that they are not directly-reachable. It is similar to other algorithms in the literature [18]. Its detailed description can be found in Figure 3.

To execute the simulation algorithms in parallel with the actions from A , we add some new (augmentation) variables to the set of variables in the states of A . Whenever a step of the simulation algorithms is executed, the state of the original variables in A remains untouched and only the new variables change their values. Whenever a process queries a local failure detector module $\mathcal{D}(p_i)$, we translate it to a query on $\text{trans}(\mathcal{D})(p_i)$. The relaying layer overlays the network with the best possible communication graph and the 3wh-layer on

top of it cuts the unidirectional edges from this graph.

Algorithm 3wh

```

1: procedure 3wh-send( $m, p_j$ )
2:   relay-send( $[1, m], p_j$ );
3:
4: procedure 3wh-recv( $m$ )
5:   relay-recv( $[l, m']$ );
6:   if ( $l = 1$ ) then
7:     relay-send( $[2, m'], sender([l, m'])$ );
8:      $m := \perp$ ;
9:   elseif ( $l = 2$ ) then
10:    relay-send( $[3, m'], sender([l, m'])$ );
11:     $m := \perp$ ;
12:   elseif ( $l = 3$ ) then
13:     $m := m'$ ;
14:   elseif  $[l, m'] = \perp$  then
15:     $m := \perp$ ;

```

Fig. 2. The Three Way Handshake Algorithm for Process p_i .

Algorithm Relay

```

1: procedure init
2:   relayed $_i := \emptyset$ ; delivered $_i := \emptyset$ ;
3:
4: procedure relay-send( $m, p_j$ )
5:   for  $k := 1$  to  $n$  do
6:     send( $[m, p_j], p_k$ );
7:     relayed $_i := relayed_i \cup \{[m, p_j]\}$ ;
8:
9: procedure relay-recv( $m$ )
10:  receive( $[m', p_k]$ );
11:  if ( $[m', p_k] = \perp$ ) then  $m := \perp$ ;
12:  elseif ( $k = i$ ) and ( $m' \notin delivered_i$ ) then
13:     $m := m'$ ; delivered $_i := delivered_i \cup \{m'\}$ ;
14:  elseif ( $k \neq i$ ) and ( $[m', p_k] \notin relayed_i$ ) then
15:    for  $l := 1$  to  $n$  do
16:      send( $[m', p_k], p_l$ );
17:      relayed $_i := relayed_i \cup \{[m', p_k]\}$ ;  $m := \perp$ ;

```

Fig. 3. The Relaying Algorithm for Process p_i .

B. The Transformation for $n > 2f$

If only less than a majority of the processes are disconnected ($n > 2f$), then we only need to adapt the problem specification to the failure patterns of the new environment. We indicate this adaptation of a problem specification with the index $p.o.$ and specify it in the following way:

$$\Sigma_{p.o.} := \{(H, \mathcal{F}) \mid \exists(H, \mathcal{F}') \in \Sigma \wedge \mathcal{F}' \leq_{om} \mathcal{F}\}$$

If we adapt consensus to omission failures, then we get Consensus $_{p.o.}$ as in Table I. The failure detector specifications

can be transformed as in Section III-A. The algorithm transformation $trans_2$ works similar as in the previous section, but we add an additional *two-way-handshake* (2wh) layer between the relaying layer and the 3wh layer. The algorithm is described in Figure 4 and is similar to an algorithm in the literature [12]. The idea of the algorithm is to broadcast every message to all other processes and to block until $f + 1$ processes have acknowledged the message. In this way, disconnected processes block forever (since they receive less than $f + 1$ acknowledgements) and connected processes can continue. Thus, we emulate a crash-stop environment.

Algorithm 2wh

```

1: procedure init
2:   received $_i := \emptyset$ ; Ack $_i := 0$ ;
3:
4: procedure 2wh-send( $m, p_j$ )
5:   relay-send( $[m, p_j, ONE], p_k$ ) to all other  $p_k$ ;
6:   Ack $_i := 1$ ;
7:   while (Ack $_i \leq f$ ) do
8:     relay-recv( $[m', p_k, num]$ );
9:     if ( $num = TWO$ ) and ( $m' = m$ )
10:      and ( $k = j$ ) then inc(Ack $_i$ );
11:     elseif ( $num = ONE$ ) then
12:       add  $[m', p_k, num]$  to received $_i$ ;
13:
14: procedure 2wh-recv( $m$ )
15:    $m := \perp$ ; relay-recv( $m'$ );
16:   if ( $m' \neq \perp$ ) then add  $m'$  to received $_i$ ;
17:   if ( $[m'', p_k, ONE] \in received_i$ ) for any  $m'', p_k$  then
18:     relay-send( $[m'', p_k, TWO]$ ,
19:               sender( $[m'', p_k, ONE]$ ));
20:   if ( $k = i$ ) then  $m := m''$ ;

```

Fig. 4. The Two Way Handshake Algorithm for Process p_i .

IV. RESULTS

In our first theorem, we show that for any algorithm A , for any failure detector \mathcal{D} , and for any problem specification Σ , $trans(A)$ using $trans(\mathcal{D})$ solves $trans(\Sigma)$ in a permanent omission environment if and only if A using \mathcal{D} solves Σ in a crash-stop environment. This theorem does not only show that our transformation works, it furthermore ensures that we do not transform to a trivial problem specification, but to an equivalent one, since we prove both directions.

Theorem 1: Let Σ be a problem specification closed under stuttering and augmentation. Then, if A is an algorithm using a failure detector \mathcal{D} and $A' = trans(A)$ is the transformation of A using $trans(\mathcal{D})$, it holds that:

$$\begin{aligned} \forall f \text{ with } 0 \leq f \leq n : \quad & (\mathcal{H}(A(\Pi), \mathcal{E}_{c.s.}^f) \subseteq \Sigma \\ \Leftrightarrow & (\mathcal{H}(A'(\Pi), \mathcal{E}_{p.o.}^f) \subseteq trans(\Sigma)) \end{aligned}$$

Proof: We divide up the proof into two parts. Let $\mathcal{S}_{c.s.} := (\mathcal{H}(A(\Pi), \mathcal{E}_{c.s.}^f)$ and $\mathcal{S}_{p.o.} := (\mathcal{H}(A'(\Pi), \mathcal{E}_{p.o.}^f)$ and assume that $A' = trans(A)$.

	Consensus	$trans(\text{Consensus})$	$\text{Consensus}_{p.o.}$
Validity:	The decided value of every process must have been proposed.	The decided value of every <i>connected</i> process must have been proposed.	The decided value of every process must have been proposed.
Non-Uniform Agreement:	No two <i>cr-correct</i> processes decide differently.	No two <i>connected</i> processes decide differently.	No two <i>connected</i> processes decide differently.
Uniform Agreement:	No two processes decide differently.	No two processes decide differently <i>before their disconnection</i> .	No two processes decide differently.
Termination:	Every <i>cr-correct</i> process eventually decides.	Every <i>connected</i> process eventually decides.	Every <i>connected</i> process eventually decides.

TABLE I
TRANSFORMATIONS OF THE CONSENSUS PROBLEM

“ \Rightarrow ”: Assume that $\mathcal{S}_{c.s.} \subseteq \Sigma$. By constructing for a given (H, \mathcal{F}) in $\mathcal{S}_{p.o.}$ a tuple (H', \mathcal{F}') in $\mathcal{S}_{c.s.}$ with $(H, \mathcal{F}) \in trans((H', \mathcal{F}'))$, we can show that $\mathcal{S}_{p.o.} \subseteq trans(\mathcal{S}_{c.s.})$ (Proposition 1). In this construction, we remove the added communication layers from H and use the properties of our two send-primitives to prove the reliability of the links in H' . We ensure “No Loss” with the relaying algorithm and “No Creation” with the three way handshake algorithm. As we know from the definition of $trans$, that $trans(\mathcal{S}_{c.s.}) \subseteq trans(\Sigma)$, we can conclude that $\mathcal{S}_{p.o.} \subseteq trans(\Sigma)$.

“ \Leftarrow ”: Assume that $\mathcal{S}_{p.o.} \subseteq trans(\Sigma)$. We construct (H', \mathcal{F}') for all (H, \mathcal{F}) in $\mathcal{S}_{c.s.}$, such that (H', \mathcal{F}') is in $\mathcal{S}_{p.o.} \subseteq trans(\Sigma)$. We can use this to prove that $\mathcal{S}_{c.s.} \subseteq \Sigma$ (Proposition 2).

Proposition 1: $\mathcal{S}_{p.o.} \subseteq trans(\mathcal{S}_{c.s.})$ *Proof:* The proposition is equivalent to

$$(H, \mathcal{F}) \in \mathcal{S}_{p.o.} \Rightarrow (H, \mathcal{F}) \in trans(\mathcal{S}_{c.s.})$$

From the definition of $trans$ follows:

$$(H, \mathcal{F}) \in \mathcal{S}_{p.o.} \Rightarrow \exists (H', \mathcal{F}') \in \mathcal{S}_{c.s.} : \\ \forall p_i \in \Pi : H'[i]/t'_{c.s.}(i) \leq_{sa} H[i]/t_{p.o.}(i) \quad (1)$$

We will in the following construct a new history H' and a failure pattern \mathcal{F}' from H and \mathcal{F} which satisfy equation (1):

- At first, we undo step 2 of the transformation and remove the variables, additional states, and events of the relaying algorithm from H . This means, that every time a relay-send or relay-receive event in H occurs, this event is substituted by an send/receive event of the underlying communication channel. We let the inserted events take place at the time when the relay events have been completed (since a process may take several steps to accomplish the relaying task). We call the intermediate history we get after this H_1 .
- Then, we undo step 1 and remove the variables, additional states, and events of the three way handshake algorithm from H_1 (in the same way as above). We call this intermediate history H_2 .

- After that, we construct \mathcal{F}' , such that $\mathcal{F}' \leq_{om} \mathcal{F}$. To build H' from H_2 , we substitute every query on a failure detector $trans(\mathcal{D})$ in H_2 with a query on \mathcal{D} in H' and remove all states and events for every process p_i that occur after the time when p_i crashes in \mathcal{F}' .

The schedule of the construction is illustrated in Figure 5. From the construction of H' and \mathcal{F}' it is clear, that $\forall p_i \in \Pi : H'[i]/t'_{c.s.}(i) \leq_{sa} H[i]/t_{p.o.}(i)$. It remains to show, that $(H', \mathcal{F}') \in \mathcal{S}_{c.s.}$. This means, that at most f processes crash in \mathcal{F}' (Lemma 1), H' is a history of $A(\Pi)$ using \mathcal{D} (Lemma 2), and all links in H' are reliable according to \mathcal{F}' (Lemma 3).

Lemma 1: At most f processes crash in \mathcal{F}' . *Proof:* Follows immediately from (c).

Lemma 2: H' is a history of $A(\Pi)$ using \mathcal{D} . *Proof:* All events and states are from $A(\Pi)$, because all additional events and states have been removed. If algorithm A makes use of a failure detector \mathcal{D} , then $trans(\mathcal{D})(\mathcal{F}) = \mathcal{D}(\mathcal{F}')$ (Since $\mathcal{F}' \leq_{om} \mathcal{F}$).

Lemma 3: All links in H' are reliable according to \mathcal{F}' . *Proof:* We have to show the three properties of reliable links, namely: No Creation (Lemma 5), No Duplication (Lemma 6), and No Loss (Lemma 7).

To prove lemma 5, we first need to show the auxiliary lemma 4:


Lemma 4: Let t_s be the time a send event from $A(p_i)$ to $A(p_j)$ in H_2 occurs, t_r be the time of the corresponding receive event in H_2 , t_j be the time when p_j becomes disconnected in \mathcal{F} , and t_i be the time when p_i become disconnected in \mathcal{F} . Then:

$$t_s \geq t_i \Rightarrow t_r \geq t_j$$

Proof: In the following, when we write $t_{nr(\mathcal{F}, p, q)}$, we mean the point in time when process p is not longer reachable from process q in \mathcal{F} (for any p, q , and \mathcal{F}).

The above lemma is equivalent to: $t_r < t_j$ implies $t_s < t_i$. At first, we observe that $t_s < t_r$. Assume $t_r < t_j$. Since $A(p_j)$ receives the message, we can conclude:

$$t_{nr(\mathcal{F}, p_j, p_i)} > t_r > t_s \quad (2)$$

Fig. 5. Construction of H'

Since the in H_2 removed 3wh-algorithm is only allowed to 3wh-deliver messages after having received a $[3, m]$ message (line 12 in Figure 2), which is only sent from a process after having on his part received a $[2, m]$ message (line 11), we are sure that after the 3wh-send event, $A(p_i)$ was able to receive the $[2, m]$ message from $A(p_j)$ and therefore:

$$t_{nr(\mathcal{F}, p_i, p_j)} > t_s \quad (3)$$

From the definition of connected follows:

$$\exists c \in \text{correct}(\mathcal{F}), t_{nr(\mathcal{F}, c, p_j)} \geq t_j > t_r > t_s \quad (4)$$

$$\exists c' \in \text{correct}(\mathcal{F}), t_{nr(\mathcal{F}, p_j, c')} \geq t_j > t_r > t_s \quad (5)$$

If we put all paths together, we have:

$$\text{with (2) \& (4) : } \exists c \in \text{correct}(\mathcal{F}), t_{nr(\mathcal{F}, c, p_i)} > t_s \quad (6)$$

$$\text{with (3) \& (5) : } \exists c' \in \text{correct}(\mathcal{F}), t_{nr(\mathcal{F}, p_i, c')} > t_s \quad (7)$$

Equations (6) and (7) imply $t_i > t_s$. ■

Lemma 5: (No Creation in H' .) For all messages m , if p_j receives m from p_i in H' , then p_i sends m to p_j in H' . *Proof:* We know, that there is no creation in H . In our construction, send events of the same layer can only decrease in the local history of crashed processes in step (c) (after the time of their crash). But since Lemma 4 shows that messages that are sent from a process that is already disconnected in \mathcal{F} (and therefore crashed in \mathcal{F}') can only be received by processes that are already disconnected too, the corresponding receive events also get lost in H' . ■

Lemma 6: (No Duplication in H' .) For all messages m : p_j receives m from p_i at most once. *Proof:* In the 3wh-algorithm, no message is delivered more than once and in the relay-algorithm, every message received is remembered in a variable $delivered_i$ (lines 12-13 in Figure 3). ■

Lemma 7: (No Loss in H' according to \mathcal{F}' .) For all messages m , if p_i sends m to p_j and p_j executes receive actions infinitely often, then p_j receives m from p_i . *Proof:* In the removed relaying algorithm, after every relay-send event, the message m is relayed by $A(p_i)$ to all other processes (lines 5-6 in Figure 3). If a connected process (in \mathcal{F}) receives such a relayed message, it checks in lines 12-13 whether it is the recipient and has not yet delivered it (and relay-delivers m in this case). Otherwise, it propagates m further to all other processes (lines 14-16).

Since p_i is at the time of the in step (a) in H_1 inserted send-event out-connected in \mathcal{F} (otherwise, p_i would have already crashed in \mathcal{F}'), there is a path of directly-reachable connected processes to a (totally) correct process in \mathcal{F} . A correct process will receive m and relay it (possibly indirectly) to $A(p_j)$, since p_j is in-connected in \mathcal{F} (because it takes infinitely many steps in (H', \mathcal{F}')). ■

Proposition 2: $\mathcal{S}_{c.s.} \subseteq \Sigma$ *Proof:* Assume $(H, \mathcal{F}) \in \mathcal{S}_{c.s.}$. We then build a new history H' from H and simulate all links according to the specification of the three-way-handshake and the relay algorithm such that $(H', \mathcal{F}) \in \text{trans}((H, \mathcal{F}))$ and $(H', \mathcal{F}) \in \mathcal{S}_{p.o.} \subseteq \text{trans}(\Sigma)$ ($\mathcal{F} \in \mathcal{E}_{c.s.}^f$ implies that $\mathcal{F} \in \mathcal{E}_{p.o.}^f$). This means, that there exists a $(H'', \mathcal{F}'') \in \Sigma$, with $(H', \mathcal{F}) \in \text{trans}((H'', \mathcal{F}''))$.

Since in both, \mathcal{F}'' and \mathcal{F} occur only crash failures, $\mathcal{F}'' = \mathcal{F}$ and therefore for all p_i , $H''[i] \leq_{sa} H'[i]$. Together with the fact that Σ is closed under stuttering and augmentation, we can conclude that $(H', \mathcal{F}) \in \Sigma$. H' and H differ only in the augmentation variables that are not relevant for the fulfillment of $\text{trans}(\Sigma)$, therefore: $(H, \mathcal{F}) \in \Sigma$. ■

Our second theorem shows, that with a majority of connected processes ($n > 2f$), trans_2 can be used to solve the adaptation of a problem to the general omission model.

Theorem 2: If A is an algorithm using a failure detector \mathcal{D} and $A' = \text{trans}_2(A)$ is the transformation of A using $\text{trans}_2(\mathcal{D})$ and Σ is closed under stuttering and augmentation, then it holds that:

$$\begin{aligned} \forall f \text{ with } f < n/2 & : (\mathcal{H}(A(\Pi), \mathcal{E}_{c.s.}^f) \subseteq \Sigma \\ & \Rightarrow (\mathcal{H}(A'(\Pi), \mathcal{E}_{p.o.}^f) \subseteq \Sigma_{p.o.} \end{aligned}$$

Proof: Let $\mathcal{S}_{c.s.} := (\mathcal{H}(A(\Pi), \mathcal{E}_{c.s.}^f)$ and $\mathcal{S}_{p.o.} := (\mathcal{H}(A'(\Pi), \mathcal{E}_{p.o.}^f)$ and assume that $A' = \text{trans}(A)$. It is sufficient to show, that

$$\begin{aligned} \forall (H, \mathcal{F}) \in \mathcal{S}_{p.o.}, \exists (H', \mathcal{F}') \in \mathcal{S}_{c.s.} : \\ (H' \leq_{sa} H) \wedge (\mathcal{F}' \leq_{om} \mathcal{F}) \end{aligned} \quad (8)$$

To show this, we construct $(H', \mathcal{F}') \in \mathcal{S}_{c.s.}$ for a given $(H, \mathcal{F}) \in \mathcal{S}_{p.o.}$ in the following way: We first remove the variables, events, and states of the relay-algorithm, then remove the same for the 2wh-algorithm, and then remove the 3wh-algorithm to get H' . \mathcal{F}' is a failure pattern, such that $\mathcal{F}' \leq_{om} \mathcal{F}$. We need to show, that (H', \mathcal{F}') fulfills the properties of equation 8. From the construction it is clear, that $H' \leq_{sa} H$ and $\mathcal{F}' \leq_{om} \mathcal{F}$. It remains to show, that $(H', \mathcal{F}') \in \mathcal{S}_{c.s.}$. This means, that at most f processes crash in \mathcal{F}' (Lemma 8), H' is a history of $A(\Pi)$ using \mathcal{D} (Lemma 9), all links are reliable in (H', \mathcal{F}') (Lemma 14), and H' and \mathcal{F}' are compatible (Lemma 13). ■

Lemma 8: At most t processes crash in \mathcal{F}' . *Proof:* Follows immediately from $\mathcal{F}' \leq_{om} \mathcal{F}$. ■

Lemma 9: H' is a history of $A(\Pi)$ using \mathcal{D} . *Proof:* All events and states are from $A(\Pi)$, because all additional events and states have been removed. If algorithm A makes use of a failure detector \mathcal{D} , then $\text{trans}(\mathcal{D})(\mathcal{F}) = \mathcal{D}(\mathcal{F}')$ (Since $\mathcal{F}' \leq_{om} \mathcal{F}$). ■

Lemma 10: Connected processes take infinitely many steps.

Proof: The only possibility for a process to block is in line 7 of the 2wh-algorithm in Figure 4. Since $n > 2f$, even after the disconnection of all f possibly faulty processes, every connected process receives acknowledgements from $n - f > f$ connected processes and therefore never blocks in line 7. ■

Lemma 11: Every process 2wh-sends at most one message after its disconnection.

Proof: If p_i is disconnected after some time, it either does not receive messages from connected processes or connected processes do not receive messages from it. If it does not receive messages from connected processes, then after a 2wh-send event, it receives at most f acknowledgements (from the disconnected ones) and therefore waits forever in line 7 of the 2wh-algorithm in Figure 4. If the connected processes do not receive messages from it and p_i 2wh-sends a message, also at most f processes will receive the ONE-message and answer with a TWO-message. Therefore, process p_i will block forever in line 7. ■

Lemma 12: The state of a process in H' does not change after its disconnection.

Proof: With the 3wh-algorithm, we can ensure that a process does not receive messages from connected processes (Lemma 4). With Lemma 11, no process sends more than one message after its disconnection (and this message is not sufficient for a 3wh). Therefore, this send event is not visible to other processes and the internal state of a disconnected process cannot be influenced after its disconnection. ■

Lemma 13: H' and \mathcal{F}' are compatible. *Proof:* We show, that every connected process takes infinitely many steps (Lemma 10), and that the state of a process after its disconnection does not change anymore in H' (Lemma 12). ■

Lemma 14: All links in H' are reliable according to \mathcal{F}' .

Proof: We have to show the three properties of reliable links, namely: No Creation (Lemma 15), No Duplication (Lemma 16), and No Loss (Lemma 17). ■

Lemma 15: No Creation in H' . *Proof:* There is no loss in H and the send events in the same layer never decrease. ■

Lemma 16: No Duplication in H' . *Proof:* In the relay-

and the 3wh-handshake algorithm, there is no duplication (Lemma 6). In the 2wh-algorithm, only one ONE message with the correct id is sent for every 2wh-send. ■

Lemma 17: No Loss in H' . *Proof:* We know from Lemma 7, that there is no loss between connected processes without the 2wh-algorithm. With Lemma 10, we know connected processes take infinitely many steps and make therefore infinitely many receive actions. It remains to show, that disconnected processes stop sending and receiving messages after their disconnection (Lemma 12). ■

1) *Weakest Failure Detectors:* A failure detector [1] is a weakest failure detector for a problem specification Σ in environment \mathcal{E} , if it is necessary and sufficient. Sufficient means, that there exists an algorithm using this failure detector that satisfies Σ in \mathcal{E} , whereas necessary means, that every other sufficient failure detector is reducible to it. A failure detector \mathcal{D} is reducible to another failure detector \mathcal{D}' , if there exists a transformation algorithm $T_{\mathcal{D} \rightarrow \mathcal{D}'}$, such that for

every tuple $(H, \mathcal{F}) \in \mathcal{H}(T_{\mathcal{D} \rightarrow \mathcal{D}'}(\Pi), \mathcal{E})$, H is equivalent to a failure detector history FDH in $\mathcal{D}'(\mathcal{F})$. We call the problem specification that arises in emulating \mathcal{D}' , $Probl(\mathcal{D}')$. In the following theorem, we show that *trans* preserves the weakest failure detector property for non-uniform¹ failure detectors.

Theorem 3: For all f with $1 \leq f \leq n$: If a non-uniform failure detector \mathcal{D} is a weakest failure detector for Σ in $\mathcal{E}_{c.s.}^f$ and Σ is closed under stuttering and augmentation, then $trans(\mathcal{D})$ is a weakest failure detector for $trans(\Sigma)$ in $\mathcal{E}_{p.o.}^f$.

Proof: If \mathcal{D} is a weakest failure detector for Σ in $\mathcal{E}_{c.s.}^f$, then $trans(\mathcal{D})$ is sufficient for $trans(\Sigma)$ in $\mathcal{E}_{p.o.}^f$ (Theorem 1). It remains to show that $trans(\mathcal{D})$ is also necessary.

Assume a failure detector \mathcal{D}' is sufficient for $trans(\Sigma)$ in $\mathcal{E}_{p.o.}^f$. Clearly, $\Sigma \subseteq trans(\Sigma)$ (since $H \leq_{sa} H$ for all H). Therefore, \mathcal{D}' is sufficient for Σ in $\mathcal{E}_{c.s.}^f$, and moreover, \mathcal{D}' is reducible to \mathcal{D} in $\mathcal{E}_{c.s.}^f$ (since \mathcal{D} is a weakest failure detector for Σ in $\mathcal{E}_{c.s.}^f$). This means that it is possible to emulate \mathcal{D} using \mathcal{D}' (i.e. a problem specification $Probl(\mathcal{D})$ that is equivalent to \mathcal{D}). If the reduction algorithm is $T_{\mathcal{D}' \rightarrow \mathcal{D}}$, then $trans(T_{\mathcal{D}' \rightarrow \mathcal{D}})$ using $trans(\mathcal{D}')$ emulates $trans(Probl(\mathcal{D}))$ in $\mathcal{E}_{p.o.}^f$ (Theorem 1) and since \mathcal{D} is non-uniform, the transformation of the problem specification, $trans(Probl(\mathcal{D}))$ is equivalent to the transformation of the failure detector $trans(\mathcal{D})$ (*trans* does not change the meaning of $Probl(\mathcal{D})$ since only the states of connected processes matter). Therefore, \mathcal{D}' is reducible to $trans(\mathcal{D})$ in $\mathcal{E}_{p.o.}^f$. ■

With Theorem 1, 2, and 3 we are able to show, the following:

Theorem 4: $trans(\Omega)$ is a weakest failure detector for uniform Consensus_{p.o.} with a majority of correct processes.

Proof: Since we know, that Ω is a weakest failure detector for non-uniform Consensus [1] and Ω is clearly non-uniform, together with Theorem 3, $trans(\Omega)$ is a weakest failure detector for non-uniform $trans(\text{Consensus})$. Since non-uniform $trans(\text{Consensus})$ is strictly weaker than uniform Consensus_{p.o.}, $trans(\Omega)$ is especially necessary for uniform Consensus_{p.o.}. To show that $trans(\Omega)$ is sufficient for uniform Consensus_{p.o.}, we can simply use Theorem 2, since we know that Ω is sufficient for uniform Consensus with a majority of correct processes. ■

V. CONCLUSION

We have given transformations for algorithms, failure detectors, and problem specifications, so crash-stop resilient algorithms can be automatically enhanced to tolerate the more severe general omission failures, highly applicable in practical settings running security problems. Furthermore we have shown that $trans(\Omega)$ is the weakest failure detector for consensus in an environment with permanent omission failures where less than half of the processes may crash. Additionally, we have proven that our transformation preserves the weakest failure detector property for all non-uniform failure detectors.

As an open problem, we think that it would be interesting to replace the requirement of a correct majority in our second transformation with a failure detector Σ [19] that will also be

¹A non-uniform failure detector \mathcal{D} outputs always the same set of histories for two failure patterns \mathcal{F} and \mathcal{F}' in which $correct(\mathcal{F}) = correct(\mathcal{F}')$ (i.e. $\mathcal{D}(\mathcal{F}) = \mathcal{D}(\mathcal{F}')$).

sufficient. Apart from that, it may be possible to give more specific transformations that are less general, but also less communication expensive than our transformation.

REFERENCES

- [1] T. D. Chandra, V. Hadzilacos, and S. Toueg, "The weakest failure detector for solving consensus," in *Proceedings of the 11th Annual ACM Symposium on Principles of Distributed Computing (PODC'92)*, M. Herlihy, Ed. Vancouver, BC, Canada: ACM Press, 1992, pp. 147–158. [Online]. Available: citeseer.ist.psu.edu/chandra96weakest.html
- [2] M. Fort, F. Freiling, L. D. Penso, Z. Benenson, and D. Kesdogan, "Trust-edpals: Secure multiparty computation implemented with smartcards," in *ESORICS '06: 11th European Symposium On Research In Computer Security*. Hamburg, Germany: Springer-Verlag, 2006, pp. 34–48.
- [3] F. Freiling, M. Herlihy, and L. D. Penso, "Optimal randomized omission-tolerant uniform consensus in message passing systems," in *9th International Conference on Principles of Distributed Systems (OPODIS)*, Dec. 2005.
- [4] G. Avoine, F. C. Gärtner, R. Guerraoui, and M. Vukolic, "Gracefully degrading fair exchange with security modules," in *The 5th European Dependable Computing Conference (EDCC)*, 2005, pp. 55–71.
- [5] V. Hadzilacos, "Issues of fault tolerance in concurrent computations (databases, reliability, transactions, agreement protocols, distributed computing)," Ph.D. dissertation, Harvard University, 1985.
- [6] K. J. Perry and S. Toueg, "Distributed agreement in the presence of processor and communication faults," *IEEE Trans. Softw. Eng.*, vol. 12, no. 3, pp. 477–482, 1986.
- [7] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.
- [8] S. Chaudhuri, "Agreement is harder than consensus: set consensus problems in totally asynchronous systems," in *Proceedings of Principles of Distributed Computing 1990*, 1990.
- [9] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of distributed consensus with one faulty process," *J. ACM*, vol. 32, no. 2, pp. 374–382, 1985.
- [10] T. D. Chandra and S. Toueg, "Unreliable failure detectors for reliable distributed systems," *Journal of the ACM*, vol. 43, no. 2, pp. 225–267, 1996. [Online]. Available: citeseer.ist.psu.edu/chandra96unreliable.html
- [11] A. Doudou, B. Garbinato, R. Guerraoui, and A. Schiper, "Muteness failure detectors: Specification and implementation," in *EDCC*, ser. Lecture Notes in Computer Science, J. Hlavicka, E. Maehle, and A. Pataricza, Eds., vol. 1667. Springer, 1999, pp. 71–87.
- [12] A. Basu, B. Charron-Bost, and S. Toueg, "Simulating reliable links with unreliable links in the presence of process crashes," in *Proceedings in the 10th International Workshop on Distributed Algorithms (WDAG96)*, 1996, pp. 105–122.
- [13] G. Neiger and S. Toueg, "Automatically increasing the fault-tolerance of distributed algorithms," *Journal of Algorithms*, vol. 11, no. 3, pp. 374–419, 1990.
- [14] C. Delporte-Gallet, H. Fauconnier, R. Guerraoui, and B. Pochon, "The perfectly-synchronised round-based model of distributed computing (to appear)," *Information & Computation*, 2007.
- [15] R. A. Bazzi and G. Neiger, "Simulating crash failures with many faulty processors (extended abstract)," in *WDAG '92: Proceedings of the 6th International Workshop on Distributed Algorithms*. London, UK: Springer-Verlag, 1992, pp. 166–184.
- [16] D. Dolev, R. Friedman, I. Keidar, and D. Malkhi, "Brief announcement: Failure detectors in omission failure environments," in *Symposium on Principles of Distributed Computing*, 1997, p. 286. [Online]. Available: citeseer.ist.psu.edu/dolev96failure.html
- [17] C. Delporte-Gallet, H. Fauconnier, and F. C. Freiling, "Revisiting failure detection and consensus in omission failure environments," in *ICTAC*, 2005, pp. 394–408.
- [18] T. K. Srikanth and S. Toueg, "Simulating authenticated broadcasts to derive simple fault-tolerant algorithms," *Distributed Computing*, vol. 2, no. 2, pp. 80–94, 1987.
- [19] C. Delporte-Gallet, H. Fauconnier, R. Guerraoui, V. Hadzilacos, P. Kouznetsov, and S. Toueg, "The weakest failure detectors to solve certain fundamental problems in distributed computing," in *PODC '04: Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*. New York, NY, USA: ACM Press, 2004, pp. 338–346.