

The Weakest Failure Detector for Eventual Consistency

Swan Dubois
Sorbonne Universités, UPMC
Univ Paris 06, CNRS, INRIA,
LIP6 UMR 7606

Rachid Guerraoui
EPFL

Petr Kuznetsov*
Télécom ParisTech

Franck Petit
Sorbonne Universités, UPMC
Univ Paris 06, CNRS, INRIA,
LIP6 UMR 7606

Pierre Sens
Sorbonne Universités, UPMC
Univ Paris 06, CNRS, INRIA,
LIP6 UMR 7606

ABSTRACT

In its classical form, a *consistent* replicated service requires all replicas to witness the same evolution of the service state. Assuming a message-passing environment with a majority of correct processes, the necessary and sufficient information about failures for implementing a general state machine replication scheme ensuring consistency is captured by the Ω failure detector.

This paper shows that in such a message-passing environment, Ω is also the weakest failure detector to implement an *eventually consistent* replicated service, where replicas are expected to agree on the evolution of the service state only after some (*a priori* unknown) time.

In fact, we show that Ω is the weakest to implement eventual consistency in *any* message-passing environment, *i.e.*, under any assumption on when and where failures might occur. Ensuring (strong) consistency in any environment requires, in addition to Ω , the quorum failure detector Σ . Our paper thus captures, for the first time, an exact computational difference between building a replicated state machine that ensures consistency and one that only ensures eventual consistency.

Categories and Subject Descriptors

C.2.4 [Computer-Communication Networks]: Distributed Systems; F.1.1 [Computation by Abstract Devices]: Models of Computation—*relations between models*

Keywords

eventual consistency, failure detectors

1. INTRODUCTION

State machine replication [21, 26] is the most studied technique to build a highly-available and consistent distributed

*The research leading to these results has received funding from the Agence Nationale de la Recherche, under grant agreement N ANR-14-CE35-0010-01, project DISCMAT.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
PODC'15, July 21–23, 2015, Donostia-San Sebastián, Spain.
Copyright © 2015 ACM 978-1-4503-3617-8/15/07 ...\$15.00.
DOI: <http://dx.doi.org/10.1145/2767386.2767404>.

service. Roughly speaking, the idea consists in replicating the service, modeled as a state machine, over several processes and ensuring that all replicas behave like one correct and available state machine, despite concurrent invocations of operations and failures of replicas. This is typically captured using the abstraction of a *total order broadcast* [3], where messages represent invocations of the service operations from clients to replicas (server processes). Assuming that the state machine is deterministic, delivering the invocations in the same total order ensures indeed that the replicas behave like a single state machine. Total order broadcast is, in turn, typically implemented by having the processes agree on which batch of messages to execute next, using the *consensus* abstraction [22, 3]. (The two abstractions, consensus and total order broadcast, were shown to be equivalent [3].)

Replicas behaving like a single one is a property generally called *consistency*. The sole purpose of the abstractions underlying the state machine replication scheme, namely consensus and total order broadcast, is precisely to ensure this consistency, while providing at the same time *availability*, namely that the replicated service does not stop responding. The inherent costs of these abstractions are sometimes considered too high, both in terms of the necessary computability assumptions about the underlying system [12, 2, 1], and the number of communication steps needed to deliver an invocation [22, 23].

An appealing approach to circumvent these costs is to trade consistency with what is sometimes called *eventual consistency* [25, 29]: namely to give up the requirement that the replicas *always* look the same, and replace it with the requirement that they only look the same *eventually*, *i.e.*, after a finite but not *a priori* bounded period of time. Basically, *eventual consistency* says that the replicas can diverge for some period, as long as this period is finite.

Many systems claim to implement general state machines that ensure eventual consistency in message-passing systems, *e.g.*, [20, 7]. But, to our knowledge, there has been no theoretical study of the exact assumptions on the information about failures underlying those implementations. This paper is the first to do so: using the formalism of failure detectors [3, 2], it addresses the question of the minimal information about failures needed to implement an eventually consistent replicated state machine.

It has been shown in [2] that, in a message-passing environment with a majority of correct processes, the weakest failure detector to implement consensus (and, thus, total order broadcast [5]) is the *eventual leader* failure detector, denoted Ω . In short, Ω outputs, at every process, a *leader*

process so that, eventually, the same correct process is considered leader by all. Ω can thus be viewed as the weakest failure detector to implement a generic replicated state machine ensuring consistency and availability in an environment with a majority of correct processes.

We show in this paper that, maybe surprisingly, the weakest failure detector to implement an *eventually consistent* replicated service in this environment (in fact, in *any* environment) is still Ω . We prove our result via an interesting generalization of the celebrated “CHT proof” by Chandra, Hadzilacos and Toueg [2]. In the CHT proof, every process periodically extracts the identifier of a process that is expected to be correct (the *leader*) from the *valencies* of an ever-growing collection of locally simulated runs. We carefully adjust the notion of valency to apply this approach to the weaker abstraction of *eventual consensus*, which we show to be necessary and sufficient to implement eventual consistency.

Our result becomes less surprising if we realize that a correct majority prevents the system from being *partitioned*, and we know that both consistency and availability cannot be achieved while tolerating partitions [1, 14, 8]. Therefore, in a system with a correct majority of processes, there is no gain in weakening consistency: (strong) consistency requires the same information about failures as eventual one. In an arbitrary environment, however, *i.e.*, under any assumptions on when and where failures may occur, the weakest failure detector for consistency is known to be $\Omega + \Sigma$, where Σ [8] returns a set of processes (called a *quorum*) so that every two such quorums intersect at any time and there is a time after which all returned quorums contain only correct processes. We show in this paper that ensuring eventual consistency does not require Σ : only Ω is needed, even if we do not assume a majority of correct processes. Therefore, Σ represents the exact difference between consistency and eventual consistency. Our result thus theoretically backs up partition-tolerance [1, 14] as one of the main motivations behind the very notion of eventual consistency.

We establish our results through the following steps:

- We give precise definitions of the notions of *eventual consensus* and *eventual total order broadcast*. We show that the two abstractions are equivalent. These underlie the intuitive notion of eventual consistency implemented in many replicated services [7, 6, 4].
- We show how to extend the celebrated CHT proof [2], initially establishing that Ω is necessary for solving consensus, to the context of eventual consensus. Through this extension, we indirectly highlight a hidden power of the technique proposed in [2] that somehow provides more than was used in the original CHT proof.
- We present an algorithm that uses Ω to implement, in any message-passing environment, an eventually consistent replicated service. The algorithm features three interesting properties: (1) An invocation can be performed after the optimal number of two communication steps, even if a majority of processes is not correct and even during periods when processes disagree on the leader, *i.e.*, partition periods;¹ (2) If Ω outputs the same leader at all processes from the very beginning, then the algorithm implements total order broadcast and hence ensures consistency; (3) *Causal* ordering is

¹Note that three communication steps are, in the worst case, necessary when strong consistency is required [23].

ensured even during periods where Ω outputs different leaders at different processes.

The rest of the paper is organized as follows. We present our system model and basic definitions in Section 2. In Section 3, we introduce abstractions for implementing eventual consistency: namely, eventual consensus and eventual total order broadcast, and we prove them to be equivalent. We show in Section 4 that the weakest failure detector for eventual consensus in any message-passing environment is Ω . We present in Section 5 our algorithm that implements eventual total order broadcast using Ω in any environment. Section 6 discusses related work, and Section 7 concludes the paper.

2. PRELIMINARIES

We adopt the classical model of distributed systems provided with the failure detector abstraction proposed in [3, 2]. In particular we employ the simplified version of the model proposed in [15, 18].

We consider a message-passing system with a set of processes $\Pi = \{p_1, p_2, \dots, p_n\}$ ($n \geq 2$). Processes execute steps of computation asynchronously, *i.e.*, there is no bound on the delay between steps. However, we assume a discrete global clock to which the processes do not have access. The range of this clock’s ticks is \mathbb{N} . Each pair of processes are connected by a reliable link.

Processes may fail by *crashing*. A *failure pattern* is a function $F : \mathbb{N} \rightarrow 2^\Pi$, where $F(t)$ is the set of processes that have crashed by time t . We assume that processes never recover from crashes, *i. e.*, $F(t) \subseteq F(t + 1)$. Let $\text{faulty}(F) = \bigcup_{t \in \mathbb{N}} F(t)$ be the set of *faulty* processes in a failure pattern F ; and $\text{correct}(F) = \Pi - \text{faulty}(F)$ be the set of *correct* processes in F . An *environment*, denoted \mathcal{E} , is a set of failure patterns.

A *failure detector history* H with range \mathcal{R} is a function $H : \Pi \times \mathbb{N} \rightarrow \mathcal{R}$, where $H(p, t)$ is interpreted as the value output by the failure detector module of process p at time t . A *failure detector* \mathcal{D} with range \mathcal{R} is a function that maps every failure pattern F to a nonempty set of failure detector histories. $\mathcal{D}(F)$ denotes the set of all possible failure detector histories that may be output by \mathcal{D} in a failure pattern F .

For example, at each process, the *leader failure detector* Ω outputs the id of a process; furthermore, if a correct process exists, then there is a time after which Ω outputs the id of the same correct process at every correct process. Another example is the *quorum failure detector* Σ , which outputs a set of processes at each process. Any two sets output at any times and by any processes intersect, and eventually every set output at any correct process consists of only correct processes.

An *algorithm* \mathcal{A} is modeled as a collection of n deterministic automata, where $\mathcal{A}(p)$ specifies the behavior of process p . Computation proceeds in *steps* of these automata. In each step, identified as a tuple (p, m, d, \mathcal{A}) , a process p atomically (1) receives a single message m (that can be the empty message λ) or accepts an *input* (from the external world), (2) queries its local failure detector module and receives a value d , (3) changes its state according to $\mathcal{A}(p)$, and (4) sends a message specified by $\mathcal{A}(p)$ for the new state to every process or produces an *output* (to the external world). Note that the use of λ ensures that a step of a process is always enabled, even if no message is sent to it.

A *configuration* of an algorithm \mathcal{A} specifies the local state of each process and the set of messages in transit. In the *initial* configuration of \mathcal{A} , no message is in transit and each process p is in the initial state of the automaton $\mathcal{A}(p)$. A

schedule S of \mathcal{A} is a finite or infinite sequence of steps of \mathcal{A} that respects $\mathcal{A}(p)$ for each p .

Following [18], we model inputs and outputs of processes using *input histories* H_I and *output histories* H_O that specify the inputs each process receives from its application and the outputs each process returns to the application over time. A run of algorithm \mathcal{A} using failure detector \mathcal{D} in environment \mathcal{E} is a tuple $R = (F, H, H_I, H_O, S, T)$, where F is a failure pattern in \mathcal{E} , H is a failure detector history in $\mathcal{D}(F)$, H_I and H_O are input and output histories of \mathcal{A} , S is a schedule of \mathcal{A} , and T is a list of increasing times in \mathbb{N} , where $T[i]$ is the time when step $S[i]$ is taken. $H \in \mathcal{D}(F)$, the failure detector values received by steps in S are consistent with H , and H_I and H_O are consistent with S . An infinite run of \mathcal{A} is *admissible* if (1) every correct process takes an infinite number of steps in S ; and (2) each message sent to a correct process is eventually received.

We then define a distributed-computing *problem*, such as consensus or total order broadcast, as a set of tuples (H_I, H_O) where H_I is an input history and H_O is an output history. An algorithm \mathcal{A} using a failure detector \mathcal{D} solves a problem P in an environment \mathcal{E} if in every admissible run of \mathcal{A} in \mathcal{E} , the input and output histories are in P . Typically, inputs and outputs represent invocations and responses of *operations* exported by the implemented abstraction. If there is an algorithm that solves P using \mathcal{D} , we sometimes, with a slight language abuse, say that \mathcal{D} *implements* P .

Consider two problems P and P' . A *transformation from P to P' in an environment \mathcal{E}* [17] is a map $T_{P \rightarrow P'}$ that, given any algorithm \mathcal{A}_P solving P in \mathcal{E} , yields an algorithm $\mathcal{A}_{P'}$ solving P' in \mathcal{E} . The transformation is *asynchronous* in the sense that \mathcal{A}_P is used as a “black box” where $\mathcal{A}_{P'}$ is obtained by feeding inputs to \mathcal{A}_P and using the returned outputs to solve P' . Hence, if P is solvable in \mathcal{E} using a failure detector \mathcal{D} , the existence of a transformation $T_{P \rightarrow P'}$ in \mathcal{E} establishes that P' is also solvable in \mathcal{E} using \mathcal{D} . If, additionally, there exists a transformation from P' to P in \mathcal{E} , we say that P and P' are *equivalent in \mathcal{E}* .

Failure detectors can be partially ordered based on their “power”: failure detector \mathcal{D} is *weaker than* failure detector \mathcal{D}' in \mathcal{E} if there is an algorithm that *emulates* the output of \mathcal{D} using \mathcal{D}' in \mathcal{E} [2, 18]. If \mathcal{D} is weaker than \mathcal{D}' , any problem that can be solved with \mathcal{D} can also be solved with \mathcal{D}' . For a problem P , \mathcal{D}^* is the *weakest* failure detector to solve P in \mathcal{E} if (a) there is an algorithm that uses \mathcal{D}^* to solve P in \mathcal{E} , and (b) \mathcal{D}^* is weaker than any failure detector \mathcal{D} that can be used to solve P in \mathcal{E} .

3. ABSTRACTIONS FOR EVENTUAL CONSISTENCY

We define two basic abstractions that capture the notion of eventual consistency: eventual total order broadcast and eventual consensus. We show that the two abstractions are equivalent: each of them can be used to implement the other.

Eventual Total Order Broadcast (ETOB).

The *total order broadcast* (TOB) abstraction [17] exports one operation *broadcast* $TOB(m)$ and maintains, at every process p_i , an output variable d_i . Let $d_i(t)$ denote the value of d_i at time t . Intuitively, $d_i(t)$ is the sequence of messages *delivered* by time t . We write $m \in d_i(t)$ if m appears in $d_i(t)$.

A process p_i *broadcasts a message m at time t* by a call to *broadcast* $TOB(m)$. We say that a process p_i *stably delivers*

a message m at time t if p_i appends m to $d_i(t)$ and m is never removed from d_i after that, *i.e.*, $m \notin d_i(t-1)$ and $\forall t' \geq t: m \in d_i(t')$. Note that if a message is delivered but not *stably* delivered by p_i at time t , it appears in $d_i(t)$ but not in $d_i(t')$ for some $t' > t$.

Assuming that broadcast messages are distinct, the TOB abstraction satisfies:

TOB-Validity If a correct process p_i broadcasts a message m at time t , then p_i eventually stably delivers m , *i.e.*, $\forall t'' \geq t: m \in d_i(t'')$ for some $t' > t$.

TOB-No-creation If $m \in d_i(t)$, then m was broadcast by some process p_j at some time $t' < t$.

TOB-No-duplication No message appears more than once in $d_i(t)$.

TOB-Agreement If a message m is stably delivered by some correct process p_i at time t , then m is eventually stably delivered by every correct process p_j .

TOB-Stability For any correct process p_i , $d_i(t_1)$ is a prefix of $d_i(t_2)$ for all $t_1, t_2 \in \mathbb{N}$, $t_1 \leq t_2$.

TOB-Total-order Let p_i and p_j be any two correct processes such that two messages m_1 and m_2 appear in $d_i(t)$ and $d_j(t)$ at time t . If m_1 appears before m_2 in $d_i(t)$, then m_1 appears before m_2 in $d_j(t)$.

We then introduce the *eventual total order broadcast* (ETOB) abstraction, which maintains the same inputs and outputs as TOB (messages are broadcast by a call to *broadcast* $ETOB(m)$) and satisfies, in every admissible run, the TOB-Validity, TOB-No-creation, TOB-No-duplication, and TOB-Agreement properties, plus the following relaxed properties for some $\tau \in \mathbb{N}$:

ETOB-Stability For any correct process p_i , $d_i(t_1)$ is a prefix of $d_i(t_2)$ for all $t_1, t_2 \in \mathbb{N}$, $\tau \leq t_1 \leq t_2$.

ETOB-Total-order Let p_i and p_j be correct processes such that messages m_1 and m_2 appear in $d_i(t)$ and $d_j(t)$ for some $t \geq \tau$. If m_1 appears before m_2 in $d_i(t)$, then m_1 appears before m_2 in $d_j(t)$.

As we show in this paper, satisfying the following optional (but useful) property in ETOB does not require more information about failures.

TOB-Causal-Order Let p_i be a correct process such that two messages m_1 and m_2 appear in $d_i(t)$ at time $t \in \mathbb{N}$. If m_2 depends causally of m_1 , then m_1 appears before m_2 in $d_i(t)$.

Here we say that a message m_2 *causally depends* on a message m_1 in a run R , and write $m_1 \rightarrow_R m_2$, if one of the following conditions holds in R : (1) a process p_i sends m_1 and then sends m_2 , (2) a process p_i receives m_1 and then sends m_2 , or (3) there exists m_3 such that $m_1 \rightarrow_R m_3$ and $m_3 \rightarrow_R m_2$.

Eventual Consensus (EC).

The *consensus* abstraction (C) [12] exports, to every process p_i , a single operation *propose* C that takes a binary argument and returns a binary response (we also say *decides*) so that the following properties are satisfied:

C-Termination Every correct process eventually returns a response to *propose* C .

C-Integrity Every process returns a response at most once.

C-Agreement No two processes return different values.

C-Validity Every value returned was previously proposed.

The *eventual consensus* (EC) abstraction exports, to every process p_i , operations $proposeEC_1, proposeEC_2, \dots$ that take binary arguments and return binary responses. Assuming that, for all $j \in \mathbb{N}$, every process invokes $proposeEC_j$ as soon as it returns a response to $proposeEC_{j-1}$, the abstraction guarantees that, in every admissible run, there exists $k \in \mathbb{N}$, such that the following properties are satisfied:

EC-Termination Every correct process eventually returns a response to $proposeEC_j$ for all $j \in \mathbb{N}$.

EC-Integrity No process responds twice to $proposeEC_j$ for all $j \in \mathbb{N}$.

EC-Validity Every value returned to $proposeEC_j$ was previously proposed to $proposeEC_j$ for all $j \in \mathbb{N}$.

EC-Agreement No two processes return different values to $proposeEC_j$ for all $j \geq k$.

It is straightforward to transform the binary version of EC into a multivalued one with unbounded set of inputs [24]. In the following, by referring to EC we mean a multivalued version of it.

Equivalence between EC and ETOB.

It is well known that, in their classical forms, the consensus and the total order broadcast abstractions are equivalent [3]. In this section, we show that a similar result holds for our eventual versions of these abstractions.

The intuition behind the transformation from EC to ETOB is the following. Each time a process p_i wants to ETOB-broadcast a message m , p_i sends m to each process. Periodically, every process p_i proposes its current sequence of messages received so far to EC. This sequence is built by concatenating the last output of EC (stored in a local variable d_i) to the batch of all messages received by the process and not yet present in d_i . The output of EC is stored in d_i , i.e., at any time, each process delivers the last sequence of messages returned by EC.

The correctness of this transformation follows from the fact that EC eventually returns consistent responses to the processes. Thus, eventually, all processes agree on the same linearly growing sequence of stably delivered messages. Furthermore, every message broadcast by a correct process eventually appears either in the delivered message sequence or in the batches of not yet delivered messages at all correct processes. Thus, by EC-Validity of EC, every message ETOB-broadcast by a correct process is eventually stored in d_i of every correct process p_i forever. By construction, no message appears in d_i twice or if it was not previously ETOB-broadcast. Therefore, the transformation satisfies the properties of ETOB.

The transformation from ETOB to EC is as follows. At each invocation of the EC primitive, the process broadcasts a message using the ETOB abstraction. This message contains the proposed value and the index of the consensus instance. As soon as a message corresponding to a given eventual consensus instance is delivered by process p_i (appears in d_i), p_i returns the value contained in the message.

Since the ETOB abstraction guarantees that every process eventually stably delivers the same sequence of messages, there exists a consensus instance after which the responses of the transformation to all alive processes are identical. Moreover, by ETOB-Validity, every message ETOB-broadcast by a correct process p_i is eventually stably delivered. Thus, every correct process eventually returns from any EC-instance

it invokes. Thus, the transformation satisfies the EC specification.

THEOREM 1. *In any environment \mathcal{E} , EC and ETOB are equivalent.*

From EC to ETOB. To prove this result, it is sufficient to provide a protocol that implements ETOB in an environment \mathcal{E} knowing that there exists a protocol that implements EC in this environment. This transformation protocol $\mathcal{T}_{EC \rightarrow ETOB}$ is stated in Algorithm 1. Now, we are going to prove that $\mathcal{T}_{EC \rightarrow ETOB}$ implements ETOB. Assume that there exists a message m broadcast by a correct process p_i at time t . As p_i is correct, every correct process receives the message $push(m)$ in a finite time. Then, m appears in the set $toDeliver$ of all correct processes in a finite time. Hence, by the termination property of EC and the construction of the function $NewBatch$, there exists ℓ such that m is included in any sequence submitted to $proposeEC_\ell$. By the EC-Validity and the EC-Termination properties, we deduce that p_i stably delivers m in a finite time, that proves that $\mathcal{T}_{EC \rightarrow ETOB}$ satisfies the TOB-Validity property. If a process p_i delivers a message m at time t , then m appears in the sequence responded by its last invocation of $proposeEC_\ell$. By construction and by the EC-Validity property, this sequence contains only messages that appear in the set $toDeliver$ of a process p_j at the time p_j invokes $proposeEC_\ell$. But this set is incrementally built at the reception of messages $push$ that contains only messages broadcast by a process. This implies that $\mathcal{T}_{EC \rightarrow ETOB}$ satisfies the TOB-No-creation. As the sequence outputted at any time by any process is the response to its last invocation of $proposeEC$ and that the sequence submitted to any invocation of this primitive contains no duplicated message (by definition of the function $NewBatch$), we can deduce from the EC-Validity property that $\mathcal{T}_{EC \rightarrow ETOB}$ satisfies the TOB-No-duplication. Assume that a correct process p_i stably delivers a message m , i.e., there exists a time after which m always appears in d_i . By the algorithm, m always appears in the response of $proposeEC$ to p_i after this time. As EC-Agreement property is eventually satisfied, we can deduce that m always appears in the response of $proposeEC$ for any correct process after some time. Thus, any correct process stably delivers m , and $\mathcal{T}_{EC \rightarrow ETOB}$ satisfies the TOB-Agreement. Let τ be the time after which the EC primitive satisfies EC-Agreement and EC-Validity. Let p_i be a correct process and $\tau \leq t_1 \leq t_2$. Let ℓ_1 (respectively ℓ_2) be the integer such that $d_i(t_1)$ (respectively $d_i(t_2)$) is the response of $proposeEC_{\ell_1}$ (respectively $proposeEC_{\ell_2}$). By construction of the protocol and the EC-Agreement and EC-Validity properties, we know that, after time τ , the response of $proposeEC_\ell$ to correct processes is a prefix of the response of $proposeEC_{\ell+1}$. As we have $\ell_1 \leq \ell_2$, we can deduce that $\mathcal{T}_{EC \rightarrow ETOB}$ satisfies the ETOB-Stability property. Let p_i and p_j be two correct processes such that two messages m_1 and m_2 appear in $d_i(t)$ and $d_j(t)$ at time $t \geq \tau$. Let ℓ be the smallest integer such that m_1 and m_2 appear in the response of $proposeEC_\ell$. By the EC-Agreement property, we know that the response of $proposeEC_\ell$ is identical for all correct processes. Then, by the ETOB-Stability property proved above, that implies that, if m_1 appears before m_2 in $d_i(t)$, then m_1 appears before m_2 in $d_j(t)$. In other words, $\mathcal{T}_{EC \rightarrow ETOB}$ satisfies the ETOB-Total-order property. In conclusion, $\mathcal{T}_{EC \rightarrow ETOB}$ satisfies the ETOB specification in an environment \mathcal{E} provided that there exists a protocol that implements EC in this environment.

Algorithm 1 $\mathcal{T}_{EC \rightarrow ETOB}$: transformation from EC to ETOB for process p_i

PROOF. Output variable:

d_i : sequence of messages of M (initially empty) outputted at any time by p_i

Internal variables:

$toDeliver_i$: set of messages of M (initially empty) containing all messages received by p_i

$count_i$: integer (initially 0) that stores the number of the last instance of consensus invoked by p_i

Messages:

$push(m)$ with m a message of M

Functions:

$Send(message)$ sends $message$ to all processes (including p_i)

$NewBatch(d_i, toDeliver_i)$ returns a sequence containing all messages from the set $toDeliver_i \setminus \{m | m \in d_i\}$

On reception of broadcastETOB(m) from the application

$Send(push(m))$

On reception of $push(m)$ from p_j

$toDeliver_i := toDeliver_i \cup \{m\}$

On reception of d as response of $proposeEC_\ell$

$d_i := d$

$count_i := count_i + 1$

$proposeEC_{count_i}(d_i.NewBatch(d_i, toDeliver_i))$

On local timeout

If $count_i = 0$ then

$count_i := 1$

$proposeEC_1(NewBatch(d_i, toDeliver_i))$

From ETOB to EC.

To prove this result, it is sufficient to provide a protocol that implements EC in an environment \mathcal{E} given a protocol that implements ETOB in this environment. This transformation protocol $\mathcal{T}_{ETOB \rightarrow EC}$ is stated in Algorithm 2. Now, we are going to prove that $\mathcal{T}_{ETOB \rightarrow EC}$ implements EC.

Let p_i be a correct process that invokes $proposeEC_\ell(v)$ with $\ell \in \mathbb{N}$. Then, by fairness and the TOB-Validity property, the construction of the protocol implies that the ETOB primitive delivers the message (ℓ, v) to p_i in a finite time. By the use of the local timeout, we know that p_i returns from $proposeEC_\ell(v)$ in a finite time, that proves that $\mathcal{T}_{ETOB \rightarrow EC}$ satisfies the EC-Termination property.

The update of the variable $count_i$ to ℓ for any process p_i that invokes $proposeEC_\ell$ and the assumptions on operations $proposeEC$ ensure us that p_i executes at most once the function $DecideEC(\ell, received_i[\Omega_i, \ell])$. Hence, $\mathcal{T}_{ETOB \rightarrow EC}$ satisfies the EC-Integrity property.

Let τ be the time after which the ETOB-Stability and the ETOB-Total-order properties are satisfied. Let k be the smallest integer such that any process that invokes $proposeEC_k$ in run r invokes it after τ .

If we assume that there exist two correct processes p_i and p_j that return different values to $proposeEC_\ell$ with $\ell \geq k$, we obtain a contradiction with the ETOB-Stability, ETOB-Total-order, or TOB-Agreement property. Indeed, if p_i returns a value after time τ , that implies that this value appears in d_i and then, by the TOB-Agreement property, this value eventually appears in d_j . If p_j returns a different value from p_i , that implies that this value is the first occurrence of a message associated to $proposeEC_\ell$ in d_j at the time of the return of $proposeEC_\ell$. After that, d_j cannot satisfy simultaneously the ETOB-Stability and the ETOB-Total-

order properties. This contradiction shows that $\mathcal{T}_{ETOB \rightarrow EC}$ satisfies the EC-Agreement property.

If we assume that there exists a process p_i that returns to $proposeEC_\ell$ with $\ell \in \mathbb{N}$ a value that was not proposed to $proposeEC_\ell$, we obtain a contradiction with the TOB-No-creation property. Indeed, the return of p_i from $proposeEC_\ell$ is chosen in d_i that contains the output of the ETOB primitive and processes broadcast only proposed values. This contradiction shows that $\mathcal{T}_{ETOB \rightarrow EC}$ satisfies the EC-Validity property.

In conclusion, $\mathcal{T}_{ETOB \rightarrow EC}$ satisfies the EC specification in an environment \mathcal{E} provided that there exists a protocol that implements ETOB in this environment. \square

Algorithm 2 $\mathcal{T}_{ETOB \rightarrow EC}$: transformation from ETOB to EC for process p_i

Internal variables:

$count_i$: integer (initially 0) that stores the number of the last instances of consensus invoked by p_i

d_i : sequence of messages (initially empty) outputted to p_i by the ETOB primitive

Functions:

$First(\ell)$: returns the value v such that (ℓ, v) is the first message of the form $(\ell, *)$ in d_i if such messages exist, \perp otherwise

$DecideEC(\ell, v)$: returns the value v as response to $proposeEC_\ell$

On invocation of $proposeEC_\ell(v)$

$count_i := \ell$

$broadcastETOB((\ell, v))$

On local time out

If $First(count_i) \neq \perp$ then

$DecideEC(count_i, First(count_i))$

4. THE WEAKEST FAILURE DETECTOR FOR EC

In this section, we show that Ω is necessary and sufficient for implementing the eventual consensus abstraction EC:

THEOREM 2. *In any environment \mathcal{E} , Ω is the weakest failure detector for EC.*

Ω is necessary for EC.

Let \mathcal{E} be any environment. We show below that Ω is weaker than any failure detector \mathcal{D} that can be used to solve EC in \mathcal{E} . Recall that implementing Ω means outputting, at every process, the identifier of a *leader* process so that eventually, the same correct leader is output permanently at all correct processes.

First, we briefly recall the arguments use by Chandra et al. [2] in the original CHT proof deriving Ω from any algorithm solving consensus (to get a more detailed survey of the proof please refer to [13, Chapter 3]). The basic observation there is that a run of any algorithm using a failure detector induces a *directed acyclic graph* (DAG). The DAG contains a sample of failure detector values output by \mathcal{D} in the current run and captures causal relations between them. Each process p_i maintains a local copy of the DAG, denoted by G_i : p_i periodically queries its failure detector module, updates G_i by connecting every vertex of the DAG with the vertex containing the returned failure-detector value with an edge, and broadcasts the DAG. An edge from vertex $[p_i, d, m]$ to

vertex $[p_j, d', m']$ is thus interpreted as “ p_i queried \mathcal{D} for the m th time and obtained value d and after that p_j queried \mathcal{D} for the m' th time and obtained value d' ”. Whenever p_i receives a DAG G_j calculated earlier by p_j , p_i merges G_i with G_j . As a result, DAGs maintained by the correct processes converge to the same infinite DAG G . The DAG G_i is then used by p_i to simulate a number of runs of the given consensus algorithm \mathcal{A} for all possible inputs to the processes. All these runs are organized in the form of a *simulation tree* Υ_i . The simulation trees Υ_i maintained by the correct processes converge to the same infinite simulation tree Υ .

The outputs produced in the simulated runs of Υ_i are then used by p_i to compute the current estimate of Ω . Every vertex σ of Υ_i is assigned a valency tag based on the decisions taken in all its *extensions* (descendants of σ in Υ_i): σ is assigned a tag $v \in \{0, 1\}$ if σ has an extension in which some process decides v . A vertex is bivalent if it is assigned both 0 and 1. It is then shown in [2] that by locating the same bivalent vertex in the limit tree Υ , the correct process can eventually extract the identifier of the same correct process. (More details can be found in [2, 13].)

We show that this method, originally designed for consensus, can be extended to eventual consensus (i.e., to the weaker EC abstraction). The extension is not trivial and requires carefully adjusting the notion of valency of a vertex in the simulation tree.

LEMMA 3. *In every environment \mathcal{E} , if a failure detector \mathcal{D} implements EC in \mathcal{E} , then Ω is weaker than \mathcal{D} in \mathcal{E} .*

PROOF. Let \mathcal{A} be any algorithm that implements EC using a failure detector \mathcal{D} in an environment \mathcal{E} . As in [2], every process p_i maintains a failure detector sample stored in DAG G_i and periodically uses G_i to simulate a set of runs of \mathcal{A} for all possible sequence of inputs of EC. The simulated runs are organized by p_i in an ever-growing *simulation tree* Υ_i . A vertex of Υ_i is the schedule of a finite run of \mathcal{A} “triggered” by a path in G_i in which every process starts with invoking $proposeEC_1(v)$, for some $v \in \{0, 1\}$, takes steps using the failure detector values stipulated by the path in G_i and, once $proposeEC_\ell(v)$ is complete, eventually invokes $proposeEC_{\ell+1}(v')$, for some $v' \in \{0, 1\}$. (For the record, we equip each vertex of Υ_i with the path in G_i used to produce it.) A vertex is connected by an edge to each one-step extension of it.²

Note that in every admissible infinite simulated run, EC-Termination, EC-Integrity and EC-Validity are satisfied and that there is $k > 0$ such that for all $\ell \geq k$, the invocations and responses of $proposeEC_\ell$ satisfy the EC-Agreement.

Since processes periodically broadcast their DAGs, the simulation tree Υ_i constructed locally by a correct process p_i converges to an infinite simulation tree Υ , in the sense that every finite subtree of Υ is eventually part of Υ_i . The infinite simulation tree Υ , starting from the initial configuration of \mathcal{A} and, in the limit, contains all possible schedules that can be triggered by the paths DAGs G_i .

Consider a vertex σ in Υ identifying a unique finite schedule of a run of \mathcal{A} using \mathcal{D} in the current failure pattern F . For $k > 0$, we say that σ is *k-enabled* if $k = 1$ or σ contains a response from $proposeEC_{k-1}$ at some process. Now

²In [2], the simulated schedules form a *simulation forest*, where a distinct simulation tree corresponds to each initial configuration encoding consensus inputs. Here we follow [18]: there is a single initial configuration and inputs are encoded in the form of input histories. As a result, we get a single simulation tree where branches depend on the parameters of $proposeEC_\ell$ calls.

Algorithm 3 Locating a bivalent vertex in Υ .

```

k := 1
σ := root of Υ
while true do
  if σ is k-bivalent then break
  σ1 := a descendant of σ in which
         EC-Agreement does not hold for proposeECk
  σ2 := a descendant of σ1 in which every correct process
         completes proposeECk and receives
         all messages sent to it in σ
  choose k' > k and σ3, a descendant of σ2, such that
         k'-tag of σ3 contains {0, 1}
  k := k'
  σ := σ3

```

we associate each vertex σ in Υ with a set of *valency tags* associated with each “consensus instance” k , called the *k-tag* of σ , as follows:

- If σ is *k-enabled* and has a descendant (in Υ) in which $proposeEC_k$ returns $x \in \{0, 1\}$, then x is added to the *k-tag* of σ .
- If σ is *k-enabled* and has a descendant in which two different values are returned by $proposeEC_k$, then \perp is added to the *k-tag* of σ .

If σ is not *k-enabled*, then its *k-tag* is empty. If the *k-tag* of σ is $\{x\}$, $x \in \{0, 1\}$, we say that σ is *(k, x)-valent* (*k-univalent*). If the *k-tag* is $\{0, 1\}$, then we say that σ is *k-bivalent*. If the *k-tag* of σ contains \perp , we say that σ is *k-invalid*.

Since \mathcal{A} ensures EC-Termination in all admissible runs extending σ , each *k-enabled* vertex σ , the *k-tag* of σ is non-empty. Moreover, EC-Termination and EC-Validity imply that a vertex in which no process has invoked $proposeEC_k$ yet has a descendant in which $proposeEC_k$ returns 0 and a descendant in which $proposeEC_k$ returns 1. Indeed, a run in which only v , $v \in \{0, 1\}$ is proposed in instance k and every correct process takes enough steps must contain v as an output. Thus:

- (*) For each vertex σ , there exists $k \in \mathbb{N}$ and σ' , a descendant of σ , such that *k-tag* of σ' contains $\{0, 1\}$.

If the “limit tree” Υ contains a *k-bivalent* vertex, we can apply the arguments of [2] to extract Ω . Now we show that such a vertex exists in Υ . Then we can simply let every process locate the “first” such vertex in its local tree Υ_i . To establish an order on the vertices, we can associate each vertex σ of Υ with the value m such that vertex $[p_i, d, m]$ of G is used to simulate the last step of σ (recall that we equip each vertex of Υ with the corresponding path). Then we order vertices of Υ in the order consistent with the growth of m . Since every vertex in G has only finitely many incoming edges, the sets of vertices having the same value of m are finite. Thus, we can break the ties in the m -based order using any deterministic procedure on these finite sets.

Eventually, by choosing the first *k-bivalent* vertex in their local trees Υ_i , the correct processes will eventually stabilize on the same *k-bivalent* vertex $\tilde{\sigma}$ in the limit tree Υ and apply the CHT extraction procedure to derive the same correct process based on *k-tags* assigned to $\tilde{\sigma}$ ’s descendants.

It remains to show that Υ indeed contains a *k-bivalent* vertex for some k . Consider the procedure described in Algorithm 3 that intends to locate such a vertex, starting with the root of the tree.

For the currently considered k -enabled vertex σ that is *not* k -bivalent (if it is k -bivalent, we are done), we use (*) to locate σ_3 , a descendant of σ , such that (1) in σ_3 , two processes return different values in $proposeEC_k$ in σ_3 , (2) in σ_3 , every correct process has completed $proposeEC_k$ and has received every message sent to it in σ , and (3) the k' -tag of σ_3 contains $\{0, 1\}$.

Thus, the procedure in Algorithm 3 either terminates by locating a k -bivalent tag and then we are done, or it never terminates. Suppose, by contradiction, that the procedure never terminates. Hence, we have an infinite admissible run of \mathcal{A} in which no agreement is provided in infinitely many instances of consensus. Indeed, in the constructed path along the tree, every correct process appears infinitely many times and receives every message sent to it. This admissible run violated the EC-Agreement property of EC—a contradiction.

Thus, the correct processes will eventually locate the same k -bivalent vertex and then, as in [2], stabilize extracting the same correct process identifier to emulate Ω . \square

Ω is sufficient for EC.

Chandra and Toueg proved that Ω is sufficient to implement the classical version of the consensus abstraction in an environment where a majority of processes are correct [3]. In this section, we extend this result to the eventual consensus abstraction for any environment.

The proposed implementation of EC is very simple. Each process has access to an Ω failure detector module. Upon each invocation of the EC primitive, a process broadcasts the proposed value (and the associated consensus index). Every process stores every received value. Each process p_i periodically checks whether it has received a value for the current consensus instance from the process that it currently believes to be the leader. If so, p_i returns this value. The correctness of this EC implementation relies on the fact that, eventually, all correct processes trust the same leader (by the definition of Ω) and then decide (return responses) consistently on the values proposed by this process.

LEMMA 4. *In every environment \mathcal{E} , EC can be implemented using Ω .*

PROOF. We propose such an implementation in Algorithm 4. Then, we prove that any admissible run r of the algorithm in any environment \mathcal{E} satisfies the EC-Termination, EC-Integrity, EC-Agreement, and EC-Validity properties.

Assume that a correct process never returns from an invocation of $proposeEC$ in r . Without loss of generality, denote by ℓ the smallest integer such that a correct process p_i never returns from the invocation of $proposeEC_\ell$. This implies that p_i always evaluates $received_i[\Omega_i, count_i]$ to \perp . We know by definition of Ω that, eventually, Ω_i always returns the same correct process p_j . Hence, by construction of ℓ , p_j returns from $proposeEC_0, \dots, proposeEC_{\ell-1}$ and then sends the message $promote(v, \ell)$ to all processes in a finite time. As p_i and p_j are correct, p_i receives this message and updates $received_i[\Omega_i, count_i]$ to v in a finite time. Therefore, the algorithm satisfies the EC-Termination property.

The update of the variable $count_i$ to ℓ for any process p_i that invokes $proposeEC_\ell$ and the assumptions on operations $proposeEC$ ensure us that p_i executes at most once the function $DecideEC(\ell, received_i[\Omega_i, \ell])$. Hence, the EC-Integrity property is satisfied.

Let τ_Ω be the time from which the local outputs of Ω are identical and constants for all correct processes in r . Let k

be the smallest integer such that any process that invokes $proposeEC_k$ in r invokes it after τ_Ω .

Let ℓ be an integer such that $\ell \geq k$. Assume that p_i and p_j are two processes that respond to $proposeEC_\ell$. Then, they respectively execute the function $DecideEC(\ell, received_i[\Omega_i, \ell])$ and $DecideEC(\ell, received_j[\Omega_j, \ell])$. By construction of k , we can deduce that $\Omega_i = \Omega_j = p_i$. That implies that p_i and p_j both received a message $promote(v, \ell)$ from p_i . As p_i sends such a message at most once, we can deduce that $received_i[p_i, \ell] = received_j[p_i, \ell]$, that proves that ensures the EC-Agreement property.

Let ℓ be an integer such that $\ell \geq k$. Assume that p_i is a process that respond to $proposeEC_\ell$. The value returned by p_i was previously received from Ω_i in a message of type $promote$. By construction of the protocol, Ω_i sends only one message of this type and this latter contains the value proposed to Ω_i , hence, the EC-Validity property is satisfied.

Thus, Algorithm 4 indeed implements EC in any environment using Ω . \square

Algorithm 4 EC using Ω : algorithm for process p_i

Local variables:

$count_i$: integer (initially 0) that stores the number of the last instances of consensus invoked by p_i
 $received_i$: two dimensional tabular that stores a value for each pair of processes/integer (initially \perp)

Functions:

$DecideEC(\ell, v)$ returns the value v as a response to $proposeEC_\ell$

Messages:

$promote(v, \ell)$ with $v \in \{0, 1\}$ and $\ell \in \mathbb{N}$

On invocation of $proposeEC_\ell(v)$

$count_i := \ell$

Send $promote(v, \ell)$ to all

On reception of $promote(v, \ell)$ from p_j

$received_i[j, \ell] := v$

On local time out

If $received_i[\Omega_i, count_i] \neq \perp$ do

$DecideEC(count_i, received_i[\Omega_i, count_i])$

5. AN EVENTUAL TOTAL ORDER BROADCAST ALGORITHM

We have shown in the previous section that Ω is the weakest failure detector for the EC abstraction (and, by Theorem 1, the ETOB abstraction) in any environment. In this section, we describe an algorithm that directly implements ETOB using Ω and which we believe is interesting in its own right.

The algorithm has three interesting properties. First, it needs only two communication steps to deliver any message when the leader does not change, whereas algorithms implementing classical TOB need at least three communication steps in this case. Second, the algorithm actually implements total order broadcast if Ω outputs the same leader at all processes from the very beginning. Third, the algorithm additionally ensures the property of TOB-Causal-Order, which does not require more information about faults.

The intuition behind this algorithm is as follows. Every process that intends to ETOB-broadcast a message sends it to all other processes. Each process p_i has access to an Ω failure detector module and maintains a DAG that stores the set of messages delivered so far together with their causal

Algorithm 5 \mathcal{ETOB} : protocol for process p_i

Output variable:

d_i : sequence of messages $m \in M$ (initially empty)
output by p_i

Internal variables:

$promote_i$: sequence of messages $m \in M$ (initially empty) promoted by p_i when $\Omega_i = p_i$
 CG_i : directed graph on messages of M (initially empty) that contains causality dependencies known by p_i

Messages:

$update(CG_i)$ with CG_i a directed graph on messages of M
 $promote(promote_i)$ with $promote_i$ a sequence of messages $m \in M$

Functions:

$UpdateCG(m, C(m))$ adds the node m and the set of edges $\{(m', m) | m' \in C(m)\}$ to CG_i
 $UnionCG(CG_j)$ replaces CG_i by the union of CG_i and CG_j
 $UpdatePromote()$ replaces $promote_i$ by one of the sequences of messages s such that $promote_i$ is a prefix of s , s contains once all messages of CG_i , and for every edge (m_1, m_2) of CG_i , m_1 appears before m_2 in s

On broadcast $\mathcal{ETOB}(m, C(m))$ **from the application**

$UpdateCG(m, C(m))$
 $Send\ update(CG_i)$ to all

On reception of $update(CG_j)$ **from** p_j

$UnionCG(CG_j)$
 $UpdatePromote()$

On reception of $promote(promote_j)$ **from** p_j

If $\Omega_i = p_j$ then
| $d_i := promote_j$

On local time out

If $\Omega_i = p_i$ then
| $Send\ promote(promote_i)$ to all

dependencies. As long as p_i considers itself the leader (its module of Ω outputs p_i), it periodically sends to all processes a sequence of messages computed from its DAG so that the sequence respects the causal order and admits the last delivered sequence as a prefix. A process that receives a sequence of messages delivers it only if it has been sent by the current leader output by Ω . The correctness of this algorithm directly follows from the properties of Ω . Indeed, once all correct processes trust the same leader, this leader promotes its own sequence of messages, which ensures the \mathcal{ETOB} specification.

The pseudocode of the algorithm is given in Algorithm 5). Below we present the proof of its correctness, including the proof that the algorithm additionally ensures TOB-Causal-Order.

LEMMA 5. *In every environment \mathcal{E} , Algorithm \mathcal{ETOB} implements \mathcal{ETOB} using Ω .*

PROOF. First, we prove that any run r of \mathcal{ETOB} in any environment \mathcal{E} satisfies the TOB-Validity, TOB-No-creation, TOB-No-duplication, and TOB-Agreement properties.

Assume that a correct process p_i broadcasts a message m at time t for a given $t \in \mathbb{N}$. We know that Ω outputs the same correct process p_j to all correct processes in a finite time. As p_j is correct, it receives the message $update(CG_i)$ from p_i (that contains m) in a finite time. Then, p_j includes m in its causality graph (by a call to $UnionCG$) and in its promotion sequence (by a call to $UpdatePromote$). As p_j never removes a message from its promotion sequence and

is outputted by Ω , p_i adopts the promotion sequence of p_j in a finite time and this sequence contains m , that proves that \mathcal{ETOB} satisfies the TOB-Validity property.

Any sequence outputted by any process is built by a call to $UpdatePromote$ by a process p_i . This function ensures that any message appearing in the computed sequence appears in the graph CG_p . This graph is built by successive calls to $UnionCG$ that ensure that the graph contains only messages received in a message of type $update$. The construction of the protocol ensures us that such messages have been broadcast by a process. Then, we can deduce that \mathcal{ETOB} satisfies the TOB-No-creation property.

Any sequence outputted by any process is built by a call to $UpdatePromote$ that ensures that any message appears only once. Then, we can deduce that \mathcal{ETOB} satisfies the TOB-No-duplication property.

Assume that a correct process p_i stably delivers a message m at time t for a given $t \in \mathbb{N}$. We know that Ω outputs the same correct process p_j to all correct processes after some finite time. Since m appears in every $d_i(t')$ such that $t' \geq t$, we derive that m appears infinitely in $promote_j$ from a given point of the run. Hence, the construction of the protocol and the correctness of p_j implies that any correct process eventually stably delivers m , and \mathcal{ETOB} satisfies the TOB-Agreement property.

We now prove that, for any environment \mathcal{E} , for any run r of \mathcal{ETOB} in \mathcal{E} , there exists a $\tau \in \mathbb{N}$ satisfying \mathcal{ETOB} -Stability, \mathcal{ETOB} -Total-order, and TOB-Causal-Order properties in r . Hence, let r be a run of \mathcal{ETOB} in an environment \mathcal{E} . Let us define:

- τ_Ω the time from which the local outputs of Ω are identical and constant for all correct processes in r ;
- Δ_c the longest communication delay between two correct processes in r ;
- Δ_t the longest local timeout for correct processes in r ;
- $\tau = \tau_\Omega + \Delta_t + \Delta_c$

Let p_i be a correct process and p_j be the correct elected by Ω after τ_Ω . Let t_1 and t_2 be two integers such that $\tau \leq t_1 \leq t_2$. As the output of Ω is stable after τ_Ω and the choice of τ ensures us that p_i receives at least one message of type $promote$ from p_j , we can deduce from the construction of the protocol that there exists $t_3 \leq t_1$ and $t_4 \leq t_2$ such that $d_i(t_1) = promote_j(t_3)$ and $d_i(t_2) = promote_j(t_4)$. But the function $UpdatePromote$ used to build $promote_j$ ensures that $promote_j(t_3)$ is a prefix of $promote_j(t_4)$. Then, \mathcal{ETOB} satisfies the \mathcal{ETOB} -Stability property after time τ .

Let p_i and p_j be two correct processes such that two messages m_1 and m_2 appear in $d_i(t)$ and $d_j(t)$ at time $t \geq \tau$. Assume that m_1 appears before m_2 in $d_i(t)$. Let p_k be the correct elected by Ω after τ_Ω . As the output of Ω is stable after τ_Ω and the choice of τ ensures us that p_i and p_j receive at least one message of type $promote$ from p_j , the construction of the protocol ensures us that we can consider t_1 and t_2 such that $d_i(t) = promote_k(t_1)$ and $d_j(t) = promote_k(t_2)$. The definition of the function $UpdatePromote$ executed by p_k allows us to deduce that either $d_i(t)$ is a prefix of $d_j(t)$ or $d_j(t)$ is a prefix of $d_i(t)$. In both cases, we obtain that m_1 appears before m_2 in $d_j(t)$, that proves that \mathcal{ETOB} satisfies the \mathcal{ETOB} -Total-order property after time τ .

Let p_i be a correct process such that two messages m_1 and m_2 appear in $d_i(t)$ at time $t \geq 0$. Assume that $m_1 \in C(m_2)$ when m_2 is broadcast. Let p_j be the process trusted by Ω_i at the time p_i adopts the sequence $d_i(t)$. If m_2 appears in

$d_i(t)$, that implies that the edge (m_1, m_2) appears in CG_j at the time p_j executes *UpdatePromote* (since p_j previously executed *UnionCG* that includes at least m and the set of edges $\{(m', m) | m' \in C(m)\}$ in CG_j). The construction of *UpdatePromote* ensures us that m_1 appears before m_2 in $d_i(t)$, that proves that \mathcal{ETOB} satisfies the TOB-Causal-Order property.

In conclusion, \mathcal{ETOB} is an implementation of ETOB assuming that processes have access to the Ω failure detector in any environment. \square

6. RELATED WORK

Modern data service providers such as Amazon’s Dynamo [7], Yahoo’s PNUTs [6] or Google Bigtable distributed storage [4] are intended to offer highly available services. They consequently replicate those services over several server processes. In order to tolerate process failures as well as partitions, they consider eventual consistency [25, 29, 28].

The term *eventual* consensus was introduced in [19]. It refers to one instance of consensus which stabilizes at the end; not multiple instances as we consider in this paper. In [9], a self-stabilizing form of consensus was proposed: assuming a self-stabilizing implementation of $\diamond S$ (also described in the paper) and executing a sequence of consensus instances, validity and agreement are eventually ensured. Their consensus abstraction is close to ours but the authors focused on the shared-memory model and did not address the question of the weakest failure detector.

In [11], the intuition behind eventual consistency was captured through the concept of eventual serializability. Two kinds of operations were defined: (1) a “stable” operation of which response needs to be totally ordered after all operations preceding it and (2) “weak” operations of which responses might not reflect all their preceding operations. Our ETOB abstraction captures consistency with respect to the “weak” operations. (Our lower bound on the necessity of Ω naturally extends to the stronger definitions.)

Our perspective on eventual consistency is closely related to the notion of *eventual linearizability* discussed recently in [27] and [16]. It is shown in [27] that the weakest failure detector to boost eventually linearizable objects to linearizable ones is $\diamond P$. We are focusing primarily on the weakest failure detector to *implement* eventual consistency, so their result is orthogonal to ours.

In [16], eventual linearizability is compared against linearizability in the context of implementing specific objects in a shared-memory context. It turns out that an eventually linearizable implementation of a *fetch-and-increment* object is as hard to achieve as a linearizable one. Our ETOB construction can be seen as an *eventually linearizable universal construction*: given any sequential object type, ETOB provides an eventually linearizable concurrent implementation of it. Brought to the message-passing environment with a correct majority, our results complement [16]: we show that in this setting, an eventually consistent replicated service (eventually linearizable object with a sequential specification) requires exactly the same information about failures as a consistent (linearizable) one.

7. CONCLUDING REMARKS

This paper defined the abstraction of eventual total order broadcast and proved its equivalence to eventual consensus: two fundamental building blocks to implement a general replicated state machine that ensures eventual consistency. We proved that the weakest failure detector to implement

these abstractions is Ω , in any message-passing environment. We could hence determine the gap between building a general replicated state machine that ensures consistency in a message-passing system and one that ensures only eventual consistency. In terms of information about failures, this gap is precisely captured by failure detector Σ [8]. In terms of time complexity, the gap is exactly one message delay: an operation on the strongly consistent replicated must, in the worst case, incur three communication steps [23], while one build using our eventually total order broadcast protocol completes an operation in the optimal number of two communication steps.

Our ETOB abstraction captures a form of eventual consistency implemented in multiple replicated services [7, 6, 4]. In addition to eventual consistency guarantees, such systems sometimes produce indications when a prefix of operations on the replicated service is *committed*, i.e., is not subject to further changes. A prefix of operations can be committed, e.g., in sufficiently long periods of synchrony, when a majority of correct processes elect the same leader and all incoming and outgoing messages of the leader to the correct majority are delivered within some fixed bound. We believe that such indications could easily be implemented, during the stable periods, on top of ETOB. Naturally, our results imply that Ω is necessary for such systems too.

Our EC abstraction assumes eventual agreement, but requires integrity and validity to be always ensured. Other definitions of eventual consensus could be considered. In particular, we have studied an eventual consensus abstraction assuming, instead of eventual agreement, *eventual integrity*, i.e., a bounded number of decisions in a given consensus instance could be revoked a finite number of times. In the companion technical report [10], we define this abstraction of eventual *irrevocable* consensus (EIC) more precisely and show that it is equivalent to our EC abstraction.

8. REFERENCES

- [1] E. A. Brewer. Towards robust distributed systems (abstract). In *Proceedings of the Nineteenth Annual ACM Symposium on Principles of Distributed Computing*, PODC '00, pages 7–, 2000.
- [2] T. D. Chandra, V. Hadzilacos, and S. Toueg. The weakest failure detector for solving consensus. *J. ACM*, 43(4):685–722, July 1996.
- [3] T. D. Chandra and S. Toueg. Unreliable failure detectors for reliable distributed systems. *J. ACM*, 43(2):225–267, Mar. 1996.
- [4] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber. Bigtable: A distributed storage system for structured data. *ACM Trans. Comput. Syst.*, 26(2):4:1–4:26, June 2008.
- [5] B. Charron-Bost and G. Tel. Approximation d’une borne inférieure répartie. Technical Report LIX/RR/94/06, Laboratoire d’Informatique LIX, École Polytechnique, Sept. 1994.
- [6] B. F. Cooper, R. Ramakrishnan, U. Srivastava, A. Silberstein, P. Bohannon, H.-A. Jacobsen, N. Puz, D. Weaver, and R. Yerneni. Pnuts: Yahoo!’s hosted data serving platform. *Proc. VLDB Endow.*, 1(2):1277–1288, Aug. 2008.
- [7] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels. Dynamo: Amazon’s highly available key-value store.

- In *Proceedings of Twenty-first ACM SIGOPS Symposium on Operating Systems Principles, SOSP '07*, pages 205–220, New York, NY, USA, 2007. ACM.
- [8] C. Delporte-Gallet, H. Fauconnier, and R. Guerraoui. Tight failure detection bounds on atomic object implementations. *J. ACM*, 57(4), 2010.
- [9] S. Dolev, R. I. Kat, and E. M. Schiller. When consensus meets self-stabilization. *Journal of Computer and System Sciences*, 76(8):884 – 900, 2010.
- [10] S. Dubois, R. Guerraoui, P. Kuznetsov, F. Petit, and P. Sens. The weakest failure detector for eventual consistency. *CoRR*, abs/X.X, 2015.
- [11] A. Fekete, D. Gupta, V. Luchangco, N. Lynch, and A. Shvartsman. Eventually-serializable data services. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing, PODC '96*, pages 300–309, New York, NY, USA, 1996. ACM.
- [12] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, Apr. 1985.
- [13] F. C. Freiling, R. Guerraoui, and P. Kuznetsov. The failure detector abstraction. *ACM Comput. Surv.*, 43(2):9:1–9:40, Feb. 2011.
- [14] S. Gilbert and N. Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, June 2002.
- [15] R. Guerraoui, V. Hadzilacos, P. Kuznetsov, and S. Toueg. The weakest failure detectors to solve quittance consensus and nonblocking atomic commit. *SIAM J. Comput.*, 41(6):1343–1379, 2012.
- [16] R. Guerraoui and E. Ruppert. A paradox of eventual linearizability in shared memory. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing, PODC '14*, pages 40–49, 2014.
- [17] V. Hadzilacos and S. Toueg. A modular approach to fault-tolerant broadcasts and related problems. Technical Report TR 94-1425, Department of Computer Science, Cornell University, May 1994.
- [18] P. Jayanti and S. Toueg. Every problem has a weakest failure detector. In *PODC*, pages 75–84, 2008.
- [19] F. Kuhn, Y. Moses, and R. Oshman. Coordinated consensus in dynamic networks. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 1–10. ACM, 2011.
- [20] A. Lakshman and P. Malik. Cassandra: A decentralized structured storage system. *SIGOPS Oper. Syst. Rev.*, 44(2):35–40, Apr. 2010.
- [21] L. Lamport. Proving the correctness of multiprocessor programs. *Transactions on software engineering*, 3(2):125–143, Mar. 1977.
- [22] L. Lamport. The Part-Time parliament. *ACM Transactions on Computer Systems*, 16(2):133–169, May 1998.
- [23] L. Lamport. Lower bounds for asynchronous consensus. *Distributed Computing*, 19(2):104–125, 2006.
- [24] A. Mostefaoui, M. Raynal, and F. Tronel. From binary consensus to multivalued consensus in asynchronous message-passing systems. *Inf. Process. Lett.*, 73(5-6):207–212, Mar. 2000.
- [25] Y. Saito and M. Shapiro. Optimistic replication. *ACM Comput. Surv.*, 37(1):42–81, Mar. 2005.
- [26] F. B. Schneider. Implementing fault-tolerant services using the state machine approach: A tutorial. *ACM Computing Surveys*, 22(4):299–319, Dec. 1990.
- [27] M. Serafini, D. Dobre, M. Majuntke, P. Bokor, and N. Suri. Eventually linearizable shared objects. In A. W. Richa and R. Guerraoui, editors, *Proceedings of the 29th Annual ACM Symposium on Principles of Distributed Computing*, pages 95–104. ACM, 2010.
- [28] A. Singh, P. Fonseca, P. Kuznetsov, R. Rodrigues, and P. Maniatis. Zoo: Eventually consistent byzantine-fault tolerance. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation, NSDI'09*, pages 169–184, Berkeley, CA, USA, 2009. USENIX Association.
- [29] W. Vogels. Eventually consistent. *Commun. ACM*, 52(1):40–44, Jan. 2009.