Revisiting Safety and Liveness in the Context of Failures^{*}

Bernadette Charron-Bost¹, Sam Toueg², and Anindya Basu³

¹ LIX, École Polytechnique, 91128 Palaiseau Cedex, France
`charron@lix.polytechnique.fr`

² DIX, École Polytechnique, 91128 Palaiseau Cedex, France

³ Bell Labs, Lucent Technologies,
600 Mountain Avenue Murray Hill, NJ 07974-0636 USA

Abstract. Safety and liveness are two fundamental concepts for proving the correctness of concurrent programs. In the context of failures, however, we observe that some properties that are commonly believed to be safety properties are actually liveness properties. In this paper, we propose refinements of the concepts of safety and liveness that avoid this counterintuitive classification.

1 Introduction

Safety and liveness are two fundamental classes of properties [Pnu77, Lam83]. Informally, a safety property asserts that something bad does not happen, and a liveness property asserts that something good eventually happens. The starting point of our work is the observation that the presence of failures can introduce some counterintuitive phenomena in this classification.

To illustrate this point, consider the well-known problem of *consensus*, and systems with process crashes. In most papers in the literature, a *correct process* is defined to be a process that takes an infinite number of steps, and the *agreement* property of consensus is defined as “no two correct processes decide differently” [FLP85, DDS87, DLS88, CHT96, AW98]. It is commonly believed that this agreement property is a safety property (e.g., [DLS88]). Upon closer look, however, it becomes clear that it is actually a liveness property. Indeed, any partial run can be extended into a run¹ that satisfies the agreement property: even if two processes have decided differently by some time t , agreement can still be satisfied if one of the two processes crashes, i.e., takes no step after time t (such a process is not correct, and so it is exempt from the agreement property). So there is a discrepancy between the intuition that agreement should be a safety property and the fact that it is actually a liveness property. Basically, this discrepancy is due to the fact that at any point, the agreement property can always be achieved thanks to future crashes.

^{*} Research partially supported by NSF grants CCR-9711403.

¹ In this paper, runs are infinite and partial runs are finite.

In this work, we propose concepts of safety and liveness that are appropriate to systems with failures in that they avoid the counterintuitive classification illustrated above. These concepts that we call *pure safety* and *pure liveness* are explained below.

First consider the concept of liveness. Roughly speaking, P is a liveness property if P says that at any point in a run, no matter what has happened up to that point, it is still possible for something good to eventually happen. In contrast, we say that P is a *pure liveness* property if P says that at any point in a run it is still possible for something good to eventually happen *and this can happen without the help of subsequent failures*. Intuitively, the motivation here is that liveness should not depend on future failures, which may or may never occur.

To illustrate the difference between liveness and pure liveness, consider the agreement property of consensus and a point in a run in which two processes have decided differently. From this point, something good may still happen to satisfy agreement — specifically, one (or both) processes may crash. Thus, as we pointed out before, agreement is a liveness property. In contrast, agreement is *not* a pure liveness property, because for something good to happen, i.e., for agreement to be satisfied, one of the two processes *must* crash at some point in the future.

To explain the concept of pure safety, we first recall the definition of safety. A partial run is *bad w.r.t some property P* , if it cannot be extended into a run that satisfies P . Property P is a *safety* property if every run that does not satisfy P has a prefix that is a bad partial run. Bad partial runs, however, are not the only undesirable ones with respect to P . Consider a partial run that cannot be extended to satisfy P without the occurrence of new failures. Intuitively, such a partial run is undesirable because failures are unpredictable, and we should not count on the occurrence of future failures to satisfy P . Formally, we say that a partial run is *undesirable w.r.t P* if every extension that satisfies P contains *additional* failures. This leads to our definition of pure safety: Property P is a *pure safety* property if every run that does not satisfy P has a prefix that is an undesirable partial run. Note that a bad partial run is also undesirable, and so pure safety is stronger than safety.

As exemplified by the agreement property of consensus, there exist liveness properties that are not pure liveness properties. Similarly, we can exhibit safety properties that are not pure safety properties. Pure safety and pure liveness are thus strict refinements of safety and liveness, respectively.

In the full paper, we first formalize the concepts of pure safety and pure liveness. To do so, we model a property as a set of runs and introduce a property transformer denoted *Pure*: roughly speaking, a property P is transformed into $Pure(P)$ by removing all the runs from P that contain an undesirable partial run. For example, for the agreement property of consensus, $Pure(agreement)$ specifies that if a process decides then its decision value is not different from the decision value of all the processes that have previously decided *and that are still alive*. In other words, $Pure(agreement)$ requires that at any time no two *alive*

processes have decided differently. Clearly, $Pure(agreement)$ is stronger than $agreement$, but weaker than $uniform\ agreement$ — a property that requires that no two processes (whether correct or faulty) ever decide differently. The property transformer $Pure$ may thus lead to some new problem specifications that are better suited to systems with failures. A property P is *pure* if it is a fixed point of the property transformer $Pure$, i.e., $Pure(P) = P$. If, in addition, P is a safety (liveness) property, then we say that P is a *pure safety* (*pure liveness*) property.

We show that pure safety and pure liveness are exempt from the counter-intuitive classification of properties illustrated above. We prove a counterpart of the well-known result of Alpern and Schneider [AS85], specifically, we show that every pure property is the conjunction of a pure safety property and a pure liveness property. We also give alternative definitions of pure safety and pure liveness properties in terms of an appropriate closure operator, which is the counterpart of the topological closure operator used to define safety and liveness in [AS85].

2 Background

In this section, we describe the model and the notation that we use. We consider a collection of processes (automata) Π and a fixed set of actions Act . Each action $a \in Act$ is signed with the identity of the process which executes a . If p executes action a , then we denote $id(a) = p$. The crash failure of $p \in \Pi$ is modeled by including in the action set of p a $crash_p$ action, the effect of which is to permanently disable all subsequent actions of p . We consider the set Act^ω of all infinite sequences on Act . The i -th action in a sequence $\sigma \in Act^\omega$ is denoted by $\sigma[i]$. The set of crash actions that occur in a sequence σ is denoted by $Crash(\sigma)$. Executions of Π correspond to a sequence in the set

$$X = \{\sigma \in Act^\omega \mid \exists i, \sigma[i] = crash_p \Rightarrow \forall j > i, id(\sigma[j]) \neq p\}.$$

Let X^* denote the set of (finite) prefixes of sequences in X . If σ and β are in X and X^* , respectively, then $\beta \prec \sigma$ means that β is a prefix of σ and $\sigma \setminus \beta$ denotes the subsequence of σ consisting of all actions that are in σ but not in β .

A process p is *correct in* $\sigma \in X$ if $crash_p$ does not occur in σ .

A *property* specifies the set of executions which it allows. This leads one to define a property formally as a predicate on X , or equivalently as a subset of X (precisely the subset of sequences satisfying the property). A sequence σ that satisfies a property P (i.e., $\sigma \in P$) is called a *trace of* P . The set of properties on X is denoted by $\mathcal{P}$. Implication, conjunction, and disjunction of properties are trivially expressed by means of set inclusion, intersection, and union, respectively.

Most of the reasoning about distributed systems has been aimed at proving two types of properties: *safety* and *liveness*. Intuitively a safety property asserts that some “bad” thing never happens. We presume that if something bad happens in a sequence, then it is as a result of some particular actions in the

sequence. This yields the following definition: a finite sequence β in X^* is *bad with respect to a property P* if β has no extension in P , i.e.,

$$\forall \sigma \succ \beta : \sigma \notin P.$$

A property P is a *safety property* if P stipulates there is no bad (finite) sequence with respect to P , namely:

$$\forall \sigma \in X : \sigma \notin P \Rightarrow \exists \beta \prec \sigma, \forall \sigma' \succ \beta : \sigma' \notin P.$$

A property P is a *liveness property* if every finite sequence in X^* has some extension in P :

$$\forall \beta \in X^*, \exists \sigma \succ \beta : \sigma \in P.$$

A liveness property is often understood as saying that some “good” thing eventually happens: no matter what happens up to some point, it is still possible for the good thing to occur at some time in the future.

There is a natural (metrical) topology on X : the basic open sets are the sets of all sequences in X which share a common prefix. With regard to this topology, the *closure of P* is the set of sequences in X such that every prefix has an extension in P :

$$\overline{P} = \{\sigma \in X \mid \forall \beta \prec \sigma, \exists \sigma' \succ \beta : \sigma' \in P\}.$$

Interestingly, safety properties are exactly the closed sets (i.e., $\overline{P} = P$), and liveness properties are exactly the dense sets (i.e., $\overline{P} = X$).

In this paper, we will illustrate our definitions and concepts with the *agreement* property of consensus. This property stipulates that no two correct processes decide differently. More precisely, let V be a fixed set of values, and assume that the set of actions of each process p contains the actions $init_p(v)$ and $decide_p(v)$, for every $v \in V$. The *agreement* property is the set of sequences σ such that

$$\left. \begin{array}{l} \forall i, \sigma[i] \notin \{crash_p, crash_q\} \\ \exists k, \sigma[k] = decide_p(v) \\ \exists l, \sigma[l] = decide_q(v') \end{array} \right\} \Rightarrow v = v'$$

As explained in the introduction, *agreement* is a liveness property.

3 Bizarre Sequences and Pure Properties

As we pointed out earlier, the *agreement* property of consensus is a liveness property. This is due to the finite sequences in X^* in which two non-crashed processes disagree: an extension of such a sequence satisfies *agreement* only if at least one of the two processes crashes. Motivated by this observation, for any property P , we define the notions of *undesirable finite sequences* and *bizarre infinite sequences with respect to P* . A finite sequence β in X^* is said to be

undesirable with respect to P if all extensions of β satisfying P contain a crash action after β :

$$\forall \sigma \in X : \beta \prec \sigma \text{ and } \sigma \in P \Rightarrow \text{Crash}(\sigma \setminus \beta) \neq \emptyset.$$

We say that an infinite sequence σ in X is *bizarre with respect to P* if it has a prefix that is undesirable with respect to P :

$$\exists \beta \prec \sigma, \forall \sigma' \in X : \beta \prec \sigma' \text{ and } \sigma' \in P \Rightarrow \text{Crash}(\sigma' \setminus \beta) \neq \emptyset.$$

The set of sequences in X that are bizarre with respect to P is denoted by $\mathcal{B}(P)$. We then define the property transformer *Pure* which removes the bizarre sequences from the traces of P . More formally,

$$\text{Pure}(P) = P \setminus \mathcal{B}(P).$$

Clearly, we have $\text{Pure}(P) \subseteq P$.

If P is a fixed point of *Pure*, i.e., $\text{Pure}(P) = P$, then P is said to be *pure*.

To illustrate this definition, we now give the pure versions of some properties. From $\text{Pure}(P) \subseteq P$, it follows that $\text{Pure}(\emptyset) = \emptyset$. Because every $\beta \in X^*$ has an extension in X with no crash action after β , X is also pure, i.e., $\text{Pure}(X) = X$. More interesting is the example of the *agreement* property of consensus: it is easy to show that $\text{Pure}(\text{agreement})$ consists of all the sequences in X for which at any point no two alive² processes have decided differently. Formally, $\text{Pure}(\text{agreement})$ is the set of sequences σ such that for any processes $p, q \in \Pi$,

$$\left. \begin{array}{l} \exists i : \sigma[i] = \text{decide}_p(v) \\ \exists j : \sigma[j] = \text{decide}_q(v') \\ i < j \text{ and } v \neq v' \end{array} \right\} \Rightarrow \exists k : i < k < j \text{ and } \sigma[k] = \text{crash}_p.$$

In order to prevent any disagreement, some works [NT90, Lyn96, CBS00] introduced a strengthening of the agreement property, called the *uniform agreement* property. More precisely, the *uniform agreement* property specifies that no two processes (whether correct or not) decide differently. The pure version of agreement only precludes disagreement among processes that are alive. The $\text{Pure}(\text{agreement})$ property is thus weaker than *uniform agreement* but stronger than *agreement*. To the best of our knowledge, this is the first time that this refinement of the *agreement* property of consensus is proposed in the literature.

In an extended version of the paper, we study the pure versions of other properties (as for example, the validity property of atomic commitment). Interestingly, the property transformer *Pure* leads to new problem specifications that are worthy studying from an algorithmic point of view.

4 Properties of the *Pure* Transformer

In the following three propositions, we state some basic properties of the property transformer *Pure*. We start with a straightforward observation:

² A process p is *alive at time k* in $\sigma \in X$ if for all $k' \leq k$, $\sigma[k'] \neq \text{crash}_p$.

Proposition 1. *Let P and Q be any two properties in $\mathcal{P}$. If $P \subseteq Q$ then $\mathcal{B}(Q) \subseteq \mathcal{B}(P)$ and $Pure(P) \subseteq Pure(Q)$.*

As an easy consequence of Proposition 1, we have the following properties of the *Pure* transformer:

Proposition 2. *Let P and Q be any two properties in $\mathcal{P}$.*

1. $Pure(P \cap Q) \subseteq Pure(P) \cap Pure(Q)$.
2. $Pure(P) \cup Pure(Q) \subseteq Pure(P \cup Q)$.

Note that part 2 of Proposition 2 implies the union (or equivalently, the disjunction) of two pure properties is pure.

Any property can be purified at most once. This is formally expressed by the following proposition:

Proposition 3. *The property transformer *Pure* is idempotent, i.e.,*

$$\forall P \in \mathcal{P} : Pure(Pure(P)) = Pure(P).$$

In other words, the pure properties (i.e., the fixed points of *Pure*) are exactly the properties in the image $Pure(\mathcal{P})$.

Proof. We show that $\mathcal{B}(Pure(P)) = \mathcal{B}(P)$. From this intermediate result, it will follow that

$$\begin{aligned} Pure(Pure(P)) &= Pure(P) \setminus \mathcal{B}(Pure(P)) \\ &= P \setminus \mathcal{B}(P) \setminus \mathcal{B}(Pure(P)) \\ &= P \setminus \mathcal{B}(P) \\ &= Pure(P). \end{aligned}$$

First, since $Pure(P) \subseteq P$, Proposition 1 implies that $\mathcal{B}(P) \subseteq \mathcal{B}(Pure(P))$. Second, we show that $\mathcal{B}(Pure(P)) \subseteq \mathcal{B}(P)$. Let σ be in $\mathcal{B}(Pure(P))$ and let β be the *shortest* prefix of σ such that

$$\forall \sigma' \succ \beta : \sigma' \in Pure(P) \Rightarrow crash(\sigma' \setminus \beta) \neq \emptyset.$$

Let $\hat{\sigma}$ be any sequence in P such that $\beta \prec \hat{\sigma}$. To complete the proof, we need to show that a crash action occurs in $\hat{\sigma} \setminus \beta$. For that, we consider two cases:

1. $\hat{\sigma} \in Pure(P)$. Since $Pure(P) \subseteq P$, it follows that $Crash(\hat{\sigma} \setminus \beta) \neq \emptyset$ as needed.
2. $\hat{\sigma} \notin Pure(P)$. Since $\hat{\sigma} \in P$, we have $\hat{\sigma} \in \mathcal{B}(P)$. Hence, there is a prefix γ of $\hat{\sigma}$ such that:

$$\forall \sigma'' \succ \gamma : \sigma'' \in P \Rightarrow Crash(\sigma'' \setminus \gamma) \neq \emptyset.$$

We consider two sub-cases:

- (a) $\beta \preceq \gamma$. Since $\hat{\sigma} \in P$, a crash action occurs in $(\hat{\sigma} \setminus \gamma)$, and *a fortiori* in $(\hat{\sigma} \setminus \beta)$.

- (b) $\gamma \prec \beta$. Let σ'' be any sequence in $Pure(P)$ such that $\gamma \prec \sigma''$. Since $Pure(P) \subseteq P$, $\sigma'' \in P$ and a crash action thus occurs in $\sigma'' \setminus \gamma$ – a contradiction with the definition of β . Therefore, this sub-case cannot occur.

As a consequence of Proposition 1 and Proposition 3, we get that every property in $\mathcal{P}$ can be given as the disjoint union of a pure property and a *completely non pure* property, i.e., a property such that its image under $Pure$ is the empty set.

Proposition 4. *For every property P in $\mathcal{P}$, there exist two properties Q and R such that:*

1. $P = Q \cup R$
2. $Pure(Q) = Q$ and $Pure(R) = \emptyset$.

Proof. By definition of $Pure(P)$, we have:

$$P = Pure(P) \cup (P \cap \mathcal{B}(P)).$$

Note that $Pure(P)$ and $P \cap \mathcal{B}(P)$ are disjoint.

Proposition 3 implies that $Pure(P)$ is pure. We claim that $Pure(P \cap \mathcal{B}(P)) = \emptyset$. By definition of the $Pure$ transformer, we have that $Pure(P \cap \mathcal{B}(P)) = (P \cap \mathcal{B}(P)) \setminus \mathcal{B}(P \cap \mathcal{B}(P))$. From Proposition 1, it follows that $\mathcal{B}(P) \subseteq \mathcal{B}(P \cap \mathcal{B}(P))$. We thus have $Pure(P \cap \mathcal{B}(P)) \subseteq (P \cap \mathcal{B}(P)) \setminus \mathcal{B}(P) = \emptyset$, i.e., $Pure(P \cap \mathcal{B}(P)) = \emptyset$ as needed.

5 The $Pure$ Transformer and the Closure Operator

For some of the developments below, it is useful to study the joint actions of the $Pure$ transformer and the topological closure operator. First, we show that any sequence that is not bizarre with respect to some property P , is necessarily a trace of $\overline{P}$.

Proposition 5. *For any property $P \in \mathcal{P}$, $X \setminus \mathcal{B}(P) \subseteq \overline{P}$.*

Proof. Let σ be any sequence in $X \setminus \mathcal{B}(P)$, i.e., for any prefix β of σ , there exists $\sigma' \in P$ that extends β and such that $Crash(\sigma' \setminus \beta) = \emptyset$. In particular, any prefix β of σ can be extended into a trace of P . This proves that σ is a trace of $\overline{P}$.

The following proposition describes the set of traces of $\overline{Pure(P)}$ in terms of the set of sequences that are bizarre with respect to P . This is a key technical result when studying the combinations of the $Pure$ transformer and the topological closure operator.

Proposition 6. *For any property $P \in \mathcal{P}$, $\overline{Pure(P)} = X \setminus \mathcal{B}(P)$.*

Proof. First, we show that if σ is a trace of $\overline{Pure(P)}$ then σ is not bizarre with respect to P . Let β be any prefix of σ ; there exists a trace of $Pure(P)$, σ' , such that $\beta \prec \sigma'$. In turn, this implies β has some extension $\hat{\sigma}$ in P such that no crash action occurs in $\hat{\sigma} \setminus \beta$. This shows that σ is not bizarre with respect to P , as needed.

We then show the converse, namely: if σ is not bizarre with respect to P then σ is a trace of $\overline{Pure(P)}$. Let β be any prefix of σ . Since $\sigma \notin \mathcal{B}(P)$, there exists an extension σ' of β in P such that no crash action occurs in $\sigma' \setminus \beta$. We now prove that σ' is a trace of $Pure(P)$. We proceed by contradiction and we assume that there exists a prefix γ of σ' such that for any $\sigma'' \in P$ extending γ , a crash action occurs in $\sigma'' \setminus \gamma$. We consider two cases:

1. $\beta \preceq \gamma$. Then $\sigma' \setminus \gamma$ is a subsequence of $\sigma' \setminus \beta$ – a contradiction to the fact that $Crash(\sigma' \setminus \beta) = \emptyset$ and $\sigma' \setminus \gamma \neq \emptyset$.
2. $\gamma \prec \beta$. Then γ is a prefix of σ . Since σ is not bizarre with respect to P , γ has an extension $\hat{\sigma}$ in P such that no crash action occurs in $\hat{\sigma} \setminus \gamma$. This contradicts the definition of γ .

Fig. 1. Grey parts represent $\overline{Pure(P)} = X \setminus \mathcal{B}(P)$.

From the two previous propositions, we get the following useful relation between $\overline{Pure(P)}$ and $Pure(\overline{P})$.

Corollary 7. For any property $P \in \mathcal{P}$, $\overline{Pure(P)} \subseteq Pure(\overline{P})$.

Proof. By Proposition 6, $\overline{Pure(P)} = X \setminus \mathcal{B}(P)$. Since $P \subseteq \overline{P}$, Proposition 1 implies that $X \setminus \mathcal{B}(P) \subseteq X \setminus \mathcal{B}(\overline{P})$. Moreover, by Proposition 5, $X \setminus \mathcal{B}(P) \subseteq \overline{P}$. It follows that

$$X \setminus \mathcal{B}(P) \subseteq (X \setminus \mathcal{B}(\overline{P})) \cap \overline{P} = Pure(\overline{P}).$$

Therefore, we have $\overline{Pure(P)} \subseteq Pure(\overline{P})$, as needed.

Inclusion in Corollary 7 is strict in general. To see that, consider the *agreement* property of consensus. As mentioned already, *agreement* is a liveness property, i.e., $\overline{\text{agreement}} = X$ and X is pure. Thus, we get $\text{Pure}(\overline{\text{agreement}}) = X$. On the other hand, by the description of $\text{Pure}(\text{agreement})$ given above, any sequence that is not a trace of $\text{Pure}(\text{agreement})$ has a prefix that cannot be extended in a trace of $\text{Pure}(\text{agreement})$. In other words, $\text{Pure}(\text{agreement})$ is a safety property, i.e., $\overline{\text{Pure}(\text{agreement})} = \text{Pure}(\text{agreement})$. This proves that $\text{Pure}(\text{agreement}) \neq \text{Pure}(\overline{\text{agreement}})$. However, it is easy to prove that if P is a pure or a safety property, then we have $\text{Pure}(P) = \text{Pure}(\overline{P})$. More precisely, we have the following two corollaries:

Corollary 8. *For any property $P \in \mathcal{P}$, if P is pure then $\overline{P}$ is also pure.*

Proof. Let P be any pure property in $\mathcal{P}$. From Corollary 7, it follows that $\overline{P} \subseteq \text{Pure}(\overline{P})$ since P is pure. By definition of the *Pure* transformer, $\text{Pure}(\overline{P}) \subseteq \overline{P}$, and so $\text{Pure}(\overline{P}) = \overline{P}$ as needed.

Corollary 9. *For any property $P \in \mathcal{P}$, we have:*

1. *If P is a safety property, then $\text{Pure}(P)$ is a safety property.*
2. *If $\text{Pure}(P)$ is a liveness property, then P is pure and P is a liveness property.*

Proof. For Part 1, let P be a safety property in $\mathcal{P}$, i.e., $\overline{P} = P$. From Corollary 7, it follows that $\overline{\text{Pure}(P)} \subseteq \text{Pure}(P)$. By definition of property closure, $\text{Pure}(P) \subseteq \overline{\text{Pure}(P)}$. Therefore, $\overline{\text{Pure}(P)} = \text{Pure}(P)$ as needed.

For Part 2, let $P \in \mathcal{P}$ such that $\text{Pure}(P) = X$. Proposition 5 implies that $\mathcal{B}(P) = \emptyset$. It follows that $\text{Pure}(P) = P$.

As explained above, the *agreement* property of consensus is a liveness property whereas $\text{Pure}(\text{agreement})$ is a safety property. This example shows that the converse of part 1 and part 2 in Corollary 9 do not hold. In other words, it may be the case that some property P is such that $\text{Pure}(P)$ is a safety property but not P . Similarly, there exist some liveness properties whose pure version is not a liveness property.

6 Most Safety Properties Are Pure

In this section, we study whether there are safety properties that are not pure. In general, the answer is yes. Indeed, let σ be a trace of X in which (at least) one crash occurs, and consider the property $P = \{\sigma\}$. Clearly, P is a safety property, i.e., $\overline{P} = P$. Let β be any prefix of σ that contains no crash action. There is exactly one extension of β that satisfies P , namely σ . By definition of β , $\sigma \setminus \beta$ contains a crash action. Therefore, $\text{Pure}(P)$ is empty, and so $\text{Pure}(P) \neq P$. In fact, as we will see below, P is a typical example of the safety properties that are not pure: a safety property that specifies the time when some process must crash is not pure.

We first formalize the notion of properties that are independent of the times at which crash failures occur. To do so, we consider the natural binary equivalence relation on X defined as follows:

$$\sigma \sim \sigma' \Leftrightarrow \text{Crash}(\sigma) = \text{Crash}(\sigma') \text{ and } \sigma \setminus \text{Crash}(\sigma) = \sigma' \setminus \text{Crash}(\sigma')$$

Informally, this means that σ and σ' have the same subsequences of non-crash actions and the same sets of crash actions. For any property $P \in \mathcal{P}$, we define property $K(P)$ by

$$K(P) = \{\sigma \in P \mid \forall \sigma' \in X, \sigma' \sim \sigma \Rightarrow \sigma' \in P\}.$$

If $K(P) = P$, then, intuitively, P is independent of the times when crash failures occur.

Proposition 10. *If P is a safety property such that $K(P) = P$, then P is pure.*

Proof. Suppose, by contradiction, that there exists $\sigma \in P \setminus \text{Pure}(P)$. Therefore σ has a prefix β such that

$$\forall \sigma' \succ \beta : \sigma' \in P \Rightarrow \text{Crash}(\sigma' \setminus \beta) \neq \emptyset.$$

In particular some crash actions occur in $\sigma \setminus \beta$. Let $(\sigma^k)_{k \in \mathbb{N}}$ be the sequence of traces such that for any index k

1. $\text{Crash}(\sigma^k) = \text{Crash}(\sigma)$;
2. $\sigma^k \setminus \text{Crash}(\sigma^k) = \sigma \setminus \text{Crash}(\sigma)$;
3. $\forall a \in \text{Crash}(\sigma)$, a does k moves to the right.

Clearly, the $(\sigma^k)_{k \in \mathbb{N}}$ sequence converges, and its limit $\tilde{\sigma}$ is an extension of β . Moreover, no crash action occurs in $\tilde{\sigma} \setminus \beta$.

From $K(P) = P$, it follows that every σ^k is in P . Since P is a safety property, $\tilde{\sigma}$ is also in P . This yields a contradiction to the fact that in each trace of P which extends β , a crash action occurs after β .

7 Pure Safety and Pure Liveness

In this section, we define two new classes of properties that are more restrictive than the classes of safety and liveness properties. In order to do so, we substitute a *pure closure* operator for the classical closure operator: a *pure safety* property will coincide with its pure closure; a *pure liveness* property will be a property whose pure closure is equal to X . We give several characterizations of pure safety properties and of pure liveness properties. We then prove a counterpart of the decomposition theorem of Alpern and Schneider [AS85], stating that every pure property is the conjunction of a pure safety property and a pure liveness property.

For any property $P \in \mathcal{P}$, we consider the property $\overline{\text{Pure}(P)}$ that we call the *pure closure* of P and denote $\tilde{P}$ (since, as shown below, the transformer $P \rightarrow \tilde{P}$

is non decreasing and idempotent, we are indeed allowed to call it a “closure”). By Proposition 6, the traces of $\tilde{P}$ are the sequences that are not bizarre with respect to P , i.e., $\tilde{P} = X \setminus \mathcal{B}(P)$. First, we prove a basic result on the pure closure operator.

Proposition 11. *For any property $P \in \mathcal{P}$, $\tilde{P}$ is both a safety property and a pure property.*

Proof. By definition, $\tilde{P}$ is a safety property. Then we claim that $\tilde{P}$ is pure. By applying Corollary 7 to $Pure(P)$, we obtain

$$\overline{Pure(Pure(\tilde{P}))} \subseteq Pure(\overline{Pure(\tilde{P})}).$$

By Proposition 3, this becomes

$$\overline{Pure(\tilde{P})} \subseteq Pure(\overline{Pure(\tilde{P})}),$$

i.e.,

$$\tilde{P} \subseteq Pure(\tilde{P}).$$

Since $Pure(\tilde{P}) \subseteq \tilde{P}$, we deduce that $Pure(\tilde{P}) = \tilde{P}$, i.e., $\tilde{P}$ is pure.

As a consequence, we can easily prove that the pure closure operator is idempotent.

Proposition 12. *The pure closure operator is idempotent, i.e.,*

$$\forall P \in \mathcal{P} : \tilde{\tilde{P}} = \tilde{P}.$$

Proof. By definition of the pure closure operator, we have $\tilde{\tilde{P}} = \overline{Pure(\tilde{P})}$. Since $\tilde{P}$ is a pure and a safety property, we get that $\tilde{\tilde{P}} = \overline{\tilde{P}}$, and then $\tilde{\tilde{P}} = \tilde{P}$.

In addition, the pure closure operator is non decreasing.

Proposition 13. *Let P and Q be any two properties in $\mathcal{P}$.*

1. *If $P \subseteq Q$, then $\tilde{P} \subseteq \tilde{Q}$.*
2. *$P \cap Q \subseteq \tilde{P} \cap \tilde{Q}$ and $\tilde{P} \cup \tilde{Q} \subseteq \widetilde{P \cup Q}$.*

Proof. Part 1 easily follows from Proposition 1 and the fact that the topological closure operator is non decreasing.

Part 2 is a straightforward consequence of Part 1.

The following proposition characterizes properties that are invariant under the pure closure operator.

Proposition 14. *For any property $P \in \mathcal{P}$, the following three assertions are equivalent:*

1. $\tilde{P} = P$.
2. P is a safety property and is pure.
3. σ is a trace of P iff σ has no undesirable prefix with respect to P , i.e.,

$$\sigma \in P \Leftrightarrow \forall \beta \prec \sigma, \exists \sigma' \succ \beta : \sigma' \in P \text{ and } \text{Crash}(\sigma' \setminus \beta) = \emptyset.$$

is called the *pure closure* of P and denoted $\tilde{P}$.

If P satisfies one of these equivalent assertions, P is said to be a *pure safety* property. Because of the last characterization, a pure safety property can be interpreted as saying that some particular “undesirable” thing never happens.

Proof. The implication (1) $\Rightarrow$ (2) is a straightforward consequence of Proposition 11. Conversely, if P is both a safety property and a pure property, then $\text{Pure}(P) = P$ and $\overline{P} = P$, and so $\tilde{P} = P$. This proves (2) $\Rightarrow$ (1).

We now show that (1) $\Leftrightarrow$ (3). By Proposition 6, (1) is equivalent to $P = X \setminus \mathcal{B}(P)$, and hence to the fact that a sequence σ is a trace of P iff σ is not bizarre with respect to P . Since a sequence is bizarre iff it has an undesirable prefix, this yields the required equivalence.

For example, the property whose set of traces is empty is trivially a pure safety property. This is also the case of property X .

At this point, it is interesting to note that Proposition 10 says that a large class of safety properties (namely the safety properties which are fixed points of the property transformer K introduced in Section 6) are indeed pure safety properties.

We now give four equivalent characterizations of the properties whose pure closure is equal to X .

Proposition 15. *For any property $P \in \mathcal{P}$, the following four assertions are equivalent:*

1. $\tilde{P} = X$.
2. P is a liveness property and is pure.
3. $\text{Pure}(P)$ is a liveness property.
4. $\forall \beta \in X^*, \exists \sigma \succ \beta : \sigma \in P \text{ and } \text{Crash}(\sigma \setminus \beta) = \emptyset$.

If P satisfies one of these equivalent assertions, P is said to be a *pure liveness* property. According to the fourth characterization, a pure liveness property can be understood as saying that there is no undesirable sequence, i.e., some particular “good” thing eventually happens without the help of crashes.

Proof. The implication (1) $\Rightarrow$ (2) follows immediately from the second point of Corollary 9. Moreover, implications (2) $\Rightarrow$ (3) and (3) $\Rightarrow$ (1) are obvious. This yields the equivalence of (1), (2) and (3).

We now show that (1) $\Rightarrow$ (4). Suppose that $\tilde{P} = X$, and let β be an arbitrary (finite) sequence of X^* . Since $\text{Pure}(P) = X$, there exists an extension σ of β that belongs to $\text{Pure}(P)$. Because $\text{Pure}(P) \subseteq P$, σ is a trace of P , and so (4) holds.

Conversely, we show that (4) $\Rightarrow$ (1). Assume that (4) holds. Let β be an arbitrary sequence of X^* . We must prove that β can be extended in a trace of $Pure(P)$. By assertion (4), β has an extension σ in P . We claim that σ is indeed in $Pure(P)$. To see this, consider any prefix γ of σ . By assertion (4), there exists an extension of γ in P in which no crash action occurs after γ . We then have $\sigma \in Pure(P)$ as needed.

We now give the counterpart of the decomposition theorem of Alpern and Schneider:

Proposition 16. *If P is a pure property, then there exist a pure safety property S and a pure liveness property L such that*

$$P = S \cap L.$$

Proof. For any property P , we trivially have:

$$P = \overline{P} \cap (X \setminus (\overline{P} \setminus P)) = \overline{P} \cap (P \cup (X \setminus \overline{P})).$$

According to [AS85] $\overline{P}$ and $P \cup (X \setminus \overline{P})$ are safety and liveness properties, respectively.

Corollary 8 states that if P is a pure property, then $\overline{P}$ is also a pure property. We now show that if P is a pure property then $P \cup (X \setminus \overline{P})$ is also a pure property. We claim that $\mathcal{B}(P \cup (X \setminus \overline{P})) = \emptyset$. Suppose for the sake of contradiction that σ is a bizarre sequence with respect to $P \cup (X \setminus \overline{P})$, i.e., $\sigma \in \mathcal{B}(P \cup (X \setminus \overline{P}))$. So σ has a finite prefix β such that

$$\forall \sigma' \succ \beta : \sigma' \in P \text{ or } \sigma' \notin \overline{P} \Rightarrow Crash(\sigma' \setminus \beta) \neq \emptyset.$$

But $Pure(X) = X$, and so there exists an extension $\tilde{\sigma}$ of β without any crash after β . By definition of β , $\tilde{\sigma} \in \overline{P} \setminus P$. In particular, β has an extension $\hat{\sigma} \in P$. On the other hand, since P is pure, some extension $\tilde{\sigma}$ of β in P is such that there is no crash action in $\tilde{\sigma}$ after β . This is a contradiction.

References

- [AS85] Bowen Alpern and Fred B. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181–185, 1985. 554, 561, 564
- [AW98] H. Attiya and J. Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics*. Mc Graw Hill, 1998. 552
- [CBS00] B. Charron-Bost and André Schiper. Uniform consensus is harder than consensus. Technical report, École Polytechnique Fédérale de Lausanne, April 2000. 556
- [CHT96] Tushar Deepak Chandra, Vassos Hadzilacos, and Sam Toueg. The weakest failure detector for solving consensus. *Journal of the ACM*, 43(4):685–722, July 1996. An extended abstract appeared in *Proceedings of the 11th ACM Symposium on Principles of Distributed Computing*, August, 1992, 147–158. 552

- [DDS87] Danny Dolev, Cynthia Dwork, and Larry Stockmeyer. On the minimal synchronism needed for distributed consensus. *Journal of the ACM*, 34(1):77–97, January 1987. 552
- [DLS88] Cynthia Dwork, Nancy A. Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM*, 35(2):288–323, April 1988. 552
- [FLP85] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32(2):374–382, April 1985. 552
- [Lam83] Leslie Lamport. What good is temporal logic? In R. E. A. Mason, editor, *Information Processing 83: proceedings of the IFIP 9th World Congress*, pages 657–668. IFIP, North-Holland, September 1983. 552
- [Lyn96] Nancy A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers, Inc., 1996. 556
- [NT90] Gil Neiger and Sam Toueg. Automatically increasing the fault-tolerance of distributed algorithms. *Journal of Algorithms*, 11(3):374–419, 1990. 556
- [Pnu77] Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th Symposium on Foundations of Computer Science*, pages 46–57. IEEE Computer Society Press, 1977. 552