

Compositionality and Observational Refinement for Linearizability with Crashes

ARTHUR OLIVEIRA VALE, ZHONGYE WANG, YIXUAN CHEN, PEIXIN YOU, and ZHONG SHAO, Yale University, USA

Crash-safety is an important property of real systems, as the main functionality of some systems is resilience to crashes. Toward a compositional verification approach for crash-safety under full-system crashes, one observes that crashes propagate instantaneously to all components across all levels of abstraction, even to unspecified components, hindering compositionality. Furthermore, in the presence of concurrency, a correctness criterion that addresses both crashes *and* concurrency proves necessary. For this, several adaptations of linearizability have been suggested, each featuring different trade-offs between complexity and expressiveness. The recently proposed compositional linearizability framework shows that to achieve compositionality with linearizability, both a locality and observational refinement property are necessary. Despite that, no linearizability criterion with crashes has been proven to support an observational refinement property.

In this paper, we define a compositional model of concurrent computation with full-system crashes. We use this model to develop a compositional theory of linearizability with crashes, which reveals a criterion, *crash-aware linearizability*, as its inherent notion of linearizability and supports both locality and observational refinement. We then show that strict linearizability and durable linearizability factor through crash-aware linearizability as two different ways of translating between concurrent computation with and without crashes, enabling simple proofs of locality and observational refinement for a generalization of these two criteria. Then, we show how the theory can be connected with a program logic for durable and crash-aware linearizability, which gives the first program logic that verifies a form of linearizability with crashes. We showcase the advantages of compositionality by verifying a library facilitating programming persistent data structures and a fragment of a transactional interface for a file system.

CCS Concepts: • **Theory of computation** → **Parallel computing models**; Denotational semantics; **Program specifications**; **Program verification**; **Abstraction**; • **Computer systems organization** → **Reliability**.

Additional Key Words and Phrases: Crash-Aware Linearizability, Strict Linearizability, Durable Linearizability, Compositional Linearizability

ACM Reference Format:

Arthur Oliveira Vale, Zhongye Wang, Yixuan Chen, Peixin You, and Zhong Shao. 2024. Compositionality and Observational Refinement for Linearizability with Crashes. *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 352 (October 2024), 115 pages. <https://doi.org/10.1145/3689792>

1 INTRODUCTION

In this paper, we develop a compositional account of linearizability under full-system crashes. By a full-system crash, we mean a crash that results in all agents of a system failing or being reset. This could result from a power outage, a user holding the power button on their computer, a fatal crash in an OS, a critical component failure, etc. By compositional, we mean that verified

Authors' address: Arthur Oliveira Vale, arthur.oliveiravale@yale.edu; Zhongye Wang, zhongye.wang@yale.edu; Yixuan Chen, yixuan.chen@yale.edu; Peixin You, peixin.you@yale.edu; Zhong Shao, zhong.shao@yale.edu, Yale University, New Haven, CT, USA.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

2475-1421/2024/10-ART352

<https://doi.org/10.1145/3689792>

components can be freely composed vertically and horizontally so that the composed system is *correct by construction*, in that no side conditions are necessary to derive its correctness from the correctness of its components. As a result, we obtain a framework for verifying large-scale crash-aware systems against linearizability. To see why compositionality is important, consider one of our main examples: the FLiT library [38].

The FLiT Library. Implementing persistent data structures, even when non-volatile memory (NVM) is available, is notoriously challenging. For instance, one of the challenges when programming with NVM is that it provides a buffered interface BCell. We can encapsulate the operations of a buffered memory cell in the following signature, where $\mathbf{1}$ stands for some singleton set (we will write $() \in \mathbf{1}$ if it is an argument, and $\text{ok} \in \mathbf{1}$ if it is a return) and Val a set of memory values:

$$\text{BCell} := \{\text{load} : \mathbf{1} \rightarrow \text{Val}, \text{store} : \text{Val} \rightarrow \mathbf{1}, \text{flush} : \mathbf{1} \rightarrow \mathbf{1}\}$$

What this signature expresses is that BCell provides three operations: $\text{load}()$, which takes unit $() \in \mathbf{1}$ as argument and returns some value in Val ; $\text{store}(v)$, which takes a value $v \in \text{Val}$ as argument, and returns the unit $\text{ok} \in \mathbf{1}$; and $\text{flush}()$, which takes unit $()$ as argument and returns a unit ok . The signature BCell provides the syntax of the operations of a buffered memory cell. It must be paired with a specification defined later, which provides the semantics of the operations. Such a specification would state that stores are not guaranteed to persist immediately; instead, they are buffered and persist only when the buffer is non-deterministically flushed or explicitly flushed by a $\text{flush}()$ invocation [33, 34]. In other words, once a crash happens, a load is only guaranteed to read a value no older than the latest flush. The explicit flush operation guarantees a buffer flush at a significant performance cost, so in practice, one would like to minimize its usage. For instance, in the trace (where $\alpha_0, \alpha_1, \alpha_2$, and α_3 are the names of the agents performing the operations):

$$\alpha_0:\text{store}(0) \cdot \alpha_0:\text{ok} \cdot \alpha_0:\text{flush}() \cdot \alpha_0:\text{ok} \cdot \alpha_1:\text{store}(1) \cdot \alpha_1:\text{ok} \cdot \alpha_2:\text{load}() \cdot \alpha_2:v \cdot \frac{1}{2} \cdot \alpha_3:\text{load}() \cdot \alpha_3:v'$$

the value v must be $v = 1$, as 1 is currently the buffered value. Meanwhile, either $v' = 0$ (the value at the latest flush), or $v' = 1$ (which could have been non-deterministically flushed from the buffer). This non-determinism of the value of a load after a crash complicates programming with NVM.

Some works attempt to facilitate programming persistent data structures by providing more robust persistent objects than those available directly from the underlying NVM, which usually only provides buffered memory cells. One such work is FLiT, a C++ library which provides a wrapper for the BCell operations. Specifically, in its essence, FLiT provides an object with signature

$$\text{FLiT} := \{\text{load} : \mathbf{1} \rightarrow \text{Val}, \text{store} : \text{Val} \rightarrow \mathbf{1}\}.$$

As is traditional in the linearizability literature, we use a set of valid concurrent traces v' to represent objects. v' may be further abstracted by providing a set v of less concurrent traces (often atomic, i.e., traces where every invocation is immediately followed by its response) with respect to which the traces in v' are linearizable¹. In the context of durable linearizability [22], v' also differs from its linearized specification v in that v' has explicit crashes while v does not. More precisely, durable linearizability requires that $\text{ops}(v')$, the crash-less specification obtained by removing all crash events from traces in v' , is linearizable (in the usual sense) w.r.t. v .

We specify the FLiT object v'_{FLiT} to be durably linearizable to v_{FLiT} , the usual crash-less atomic memory cell. This should be understood as stating that the FLiT operations are persistent in v'_{FLiT} , meaning that a load after a crash does read the most recently written value, up to happens-before reordering. FLiT's implementation M_{FLiT} , which runs on top of a buffered memory cell object v'_{BCell} and of a volatile counter object v'_{Counter} , does this by: (1) always flushing stores; (2) using the counter to keep track of when flushes are necessary; (3) having loads only flush when the counter marks that a flush is necessary. The counter specified by v'_{Counter} is volatile in that it lives in volatile

¹We take the convention that a primed specification is a concrete specification, and the un-primed an abstract specification

memory, so after a crash, a new instance is created with the initial value of 0. The code M_{FLiT} for our simplified formulation of FLiT is found below in Fig. 1.

For instance, a buffered memory cell allows for the following trace:

$$\alpha_0:\text{store}(1) \cdot \alpha_1:\text{load}() \cdot \alpha_1:1 \cdot \zeta \cdot \alpha_2:\text{load}() \cdot \alpha_2:v$$

where either $v = 0$ (when the buffer containing 1 has not flushed before the crash) or $v = 1$ (when the buffer is flushed before the crash). If $v = 0$, the trace is not durably linearizable to the usual memory cell specification because a 0 is read after 1 is read with no $\text{store}(0)$ to justify it.

Meanwhile, when using FLiT, the call to $\text{store}(1)$ must execute at least up to the $B.\text{store}(1)$ invocation (as $\text{load}()$ manages to read 1). This means that the call to $\text{store}(1)$ will have executed $C.\text{inc}()$. Assuming it only executes up to receiving the response $B.\text{ok}$ to its $B.\text{store}(1)$ call (otherwise, it executes a flush). The $\alpha_1:\text{load}()$ call will execute to completion, so it will call $B.\text{load}()$ and receive $B.1$ as response. Then, it will read 1 from $C.\text{get}()$, and will execute $B.\text{flush}()$ before returning 1. Hence, when the crash ζ happens, the buffered memory cell has been flushed, guaranteeing that any $\text{load}()$ calls after the crash will read 1. Therefore, calling the memory operations using FLiT guarantees that $v = 1$.

The FLiT paper claims that: “Using the library’s default mode makes any linearizable data structure durable [...]”, which they do not prove. In fact, it is challenging to state this theorem without a compositional model of crash-aware computation, as it concerns discussing the composition of arbitrary clients with FLiT. In addition, even if such a compositional model were available, it must provide good support for durable linearizability and be closely connected with a concurrent compositional model without crashes, also providing good support for usual linearizability [20]. The reason for this is that this statement relates an implementation that assumes the usual concurrent memory and implements a linearizable object, with an implementation that runs on top of the crash-aware FLiT library and implements a durably linearizable object. No framework for verification of concurrent systems with crashes allows for the correctness of FLiT to be stated in full formality, much less for it to be proved and used to build provably correct durable components using a crash-less component (i.e., one whose specification does not involve crashes) which has been previously verified against a linearizability specification.

Using our compositional account of linearizability with crashes, we prove the following FLiT correctness theorem (v'_{Cell} is any crash-less object Herlihy-Wing linearizable to v_{FLiT}).

PROPOSITION 1.1 (FLiT CORRECTNESS). *For any object signature E , writing $v'_{\text{Mem}} := \otimes_{i \in I} v'_{\text{Cell}}$ for the horizontal composition of several memory cells, if $v'_{\text{Mem}}; M$ is an object linearizable to v_E then, writing $v'_{\text{BMem}} := \otimes_{i \in I} v'_{\text{BCell}}$, it follows that $v'_{\text{BMem}}; M_{\text{FLiT}}; \text{vol}(M)$ is durably linearizable to v_E .*

The $-\otimes-$ operation stands for horizontal composition, which composes two objects into a single object, allowing operations from both components to be issued by a client. Therefore, v'_{Mem} defines a memory array. M is code implementing a new object with signature E using the memory array v'_{Mem} . The $-\text{;}-$ stands for vertical composition, so that $v'_{\text{Mem}}; M$ stands for the object obtained by running the implementation M on top of the memory array. Similarly, v'_{BMem} is a buffered memory array. $\text{vol}(M)$ adds crash semantics to M by running it in each epoch (the period in between crashes), so that $v'_{\text{BMem}}; M_{\text{FLiT}}; \text{vol}(M)$ is the object obtained by running M on top of the FLiT wrapper M_{FLiT} around the buffered memory array.

```

import B: BCell
import C: Counter

load()
{
  v ← B.load();
  if(C.get() != 0)
  { B.flush(); }
  return v;
}

store(v)
{
  C.inc();
  B.store(v);
  B.flush();
  C.dec();
  return;
}

```

Fig. 1. FLiT Memory Cell Implementation M_{FLiT}

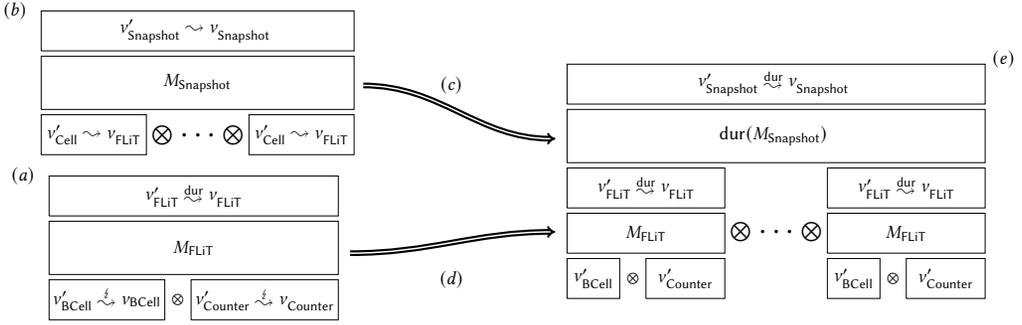


Fig. 2. (a) Using our program logic for durable linearizability, we verify the FLiT implementation; (b) Using the program logic for compositional linearizability we verify the crash-less snapshot object; (c) Using the FLiT correctness theorem, we lift the crash-less snapshot object into durably linearizable snapshot object running on top of a FLiT array; (d) Using vertical and horizontal composition, we obtain (e) a durably linearizable snapshot object running on top of an array of buffered memory cells and volatile counters.

We prove this by developing a compositional theory of durable linearizability which supports both locality and observational refinement (where we write $v' \overset{dur}{\rightsquigarrow} v$ for “ v' is durably linearizable to v ”), proving locality and observational refinement properties for durable linearizability, and then introducing a program logic for verifying individual components to be durably linearizable. Using our program logic, we show that (depicted diagrammatically in Fig. 2 (a)):

PROPOSITION 1.2. $(v'_{Counter} \otimes v'_{BCell}); M_{FLiT}$ is durably linearizable with respect to v_{FLiT} .

$(v'_{Counter} \otimes v'_{BCell}); M_{FLiT}$ is the object obtained by running the code in Fig. 1 on top of the volatile counter $v'_{Counter}$ and the buffered memory cell v'_{BCell} . By using an observational refinement property for *crash-aware linearizability*, a novel linearizability criterion we introduce, we prove this using instead the linearized specifications for the counter and the buffered memory cell, greatly simplifying its proof by only considering atomic traces. By verifying that M_{FLiT} is durably linearizable, we can use locality (Prop. 1.4) and observational refinement (Prop. 1.3) to prove FLiT’s correctness.

PROPOSITION 1.3 (OBSERVATIONAL REFINEMENT). An object $v'_A : A$ is durably linearizable to v_A if and only if whenever an implementation M implements a concurrent object linearizable to v_B using v_A , $vol(M)$ implements an object durably linearizable to v_B using v'_A .

PROPOSITION 1.4 (LOCALITY). For $v'_A : A, v'_B : B$ and $v_A : A, v_B : B$:

$$v'_A \overset{dur}{\rightsquigarrow} v_A \text{ and } v'_B \overset{dur}{\rightsquigarrow} v_B \text{ if and only if } v'_A \otimes v'_B \overset{dur}{\rightsquigarrow} v_A \otimes v_B$$

While the original paper on durable linearizability claims it satisfies locality, it does not do so by formalizing horizontal composition. Meanwhile, our locality statement is directly formulated within our compositional model of computation with crashes, which is defined independent of any notion of linearizability. This makes our locality theorem much stronger as it interacts well with refinement and vertical composition. Observational refinement, however, has never been shown for any linearizability criteria with crashes. Our program logic is the first to verify any linearizability criteria with crashes. Moreover, as it is necessary to verify the FLiT library, our program logic can reason about external linearization points and helpings [26], even across crashes.

We further showcase the benefits of compositionality and of the FLiT correctness theorem by showing that we can lift a crash-less interval-sequential linearizable snapshot object [7] into a durable one (Fig. 2 (b)). We do this by first verifying the write-snapshot implementation $M_{Snapshot}$

from Borowsky and Gafni [6] using the program logic of Oliveira Vale et al. [31, 32]. The implementation uses a (crash-less) memory cell object v'_{Cell} (which is Herlihy-Wing linearizable to v_{FLIT}) to implement an object with interface Snapshot with a single operation `write_snapshot`:

$$\text{Snapshot} := \{\text{write_snapshot} : \text{Val} \rightarrow \mathcal{P}(\text{Val})\}$$

The operation `write_snapshot` writes the current value to the memory and returns a set of values that have been written to the object before. The implementation M_{Snapshot} uses one memory cell per agent $\alpha \in S$ in the snapshot system to implement the Snapshot object.

Using the soundness theorem for their program logic [31], we obtain a crash-less interval-sequential linearizable object. Because we formally connect our model and durable linearizability definition to their model, we can then use the FLIT correctness theorem to obtain that the write-snapshot object is interval-sequential *durably* linearizable in the model with crashes. Note that this also showcases that our linearizability criteria and program logic are all generalized to handle interval-sequential objects. We display this setup in Fig. 2 (e).

While durable linearizability is a good criterion for specifying persistent objects, it is inept at expressing objects with less persistent behaviors, such as volatile objects, buffered objects, or objects with hybrid crash behaviors (e.g., horizontal compositions of objects with different persistency guarantees). Therefore, we use the methodology of compositional linearizability [31] to derive the inherent notion of linearizability to our compositional model, which we call *crash-aware linearizability*. We show that this criterion, though simple, is novel to our work and satisfies locality and observational refinement. Then, we show that durable linearizability and strict linearizability factor through crash-aware linearizability as different ways to translate crash-aware linearizable objects to the crash-less model from compositional linearizability. We showcase that crash-aware linearizability is a robust verification criterion by verifying a fragment of a transactional file system interface featuring recovery and objects with many different persistency guarantees.

Summary of Main Contributions.

- A compositional model of concurrent computation with crashes directly connected to the model of crash-less computation used in the compositional linearizability paper [31].
- A novel linearizability criterion, which we call crash-aware linearizability, is apt for specifying objects with a variety of crash behaviors.
- Compositional formulations of strict and durable linearizability, in particular, generalizing them away from atomic specifications.
- Proofs of locality, formulated for the first time in a compositional style, for crash-aware, strict, and durable linearizability.
- The first proofs of observational refinement properties for any linearizability criterion with crashes, which we show for crash-aware, strict, and durable linearizability.
- Two variations of a program logic for showing linearizability of crash-aware components: one for crash-aware linearizability and the other for durable linearizability. This makes for the first program logic that can prove linearizability specifications for components with crashes.
- A proof of correctness for FLIT and a proof that the snapshot object of Borowsky and Gafni [6] is interval-sequential linearizable, yielding a verified durable interval-sequential snapshot object using the FLIT correctness theorem.
- A proof of correctness, against crash-aware linearizability, of a simplified file API, involving objects with a variety of crash-behaviors and a few layers to exemplify compositionality under heterogenous crash-behaviors.

We present a reduced treatment of our results, which emphasizes the main points and omits all proofs. A full account of our results may be found in our extensive appendix.

2 THREE LINEARIZABILITY CRITERIA UNDER CRASHES

We start the technical core of our paper by defining and contrasting three different linearizability criteria under crashes: *crash-aware linearizability*, *strict linearizability* and *durable linearizability*. We assume a crash model with full-system crashes, that is, a crash event crashes all agents in the system. This is appropriate for, for example, a multicore machine but not for a distributed system, which requires individual crashes for each node. It serves, however, as a crucial stepping stone toward a realistic compositional modeling of distributed systems with crashes, as each node is often a multi-threaded system over a multicore machine. We define the criteria formally but omit many technical details of the compositional model, which we explain later in §3.

2.1 Preliminaries

Our model is parametrized by a set Υ of agent names $\alpha \in \Upsilon$. Events look like $\alpha:m$ denoting that agent α performs an invocation or response m . If M denotes the given set of events then $s \in M^*$ is said to be a *crash-less* well-formed trace if its projection $\pi_\alpha(s)$ to only events performed by α alternates between invocations and responses, and denote the set of all such traces by $\mathbb{P}_M^{\text{conc}}$.

We denote a crash event by \downarrow . We say a trace $s \in (M + \downarrow)^*$ is a well-formed *crash-aware* trace if it is of the form $s_1 \cdot \downarrow \cdot s_2 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_n$ where each $s_i \in \mathbb{P}_M^{\text{conc}}$. Given this decomposition, we define the number of epochs $\|s\|$ of s to be $\|s\| := n$. The trace s_i is called the i -th epoch of s and denoted by $\text{epo}_i(s) := s_i$. We denote the set of all well-formed crash-aware traces over M by \mathbb{P}_M^\downarrow .

As usual with linearizability, a specification is a non-empty, prefix-closed set of well-formed traces. If the specification ν only has crash-less traces, i.e. $\nu \subseteq \mathbb{P}_M^{\text{conc}}$, we call it a *crash-less specification*, and if it has crash-aware traces, i.e. $\nu \subseteq \mathbb{P}_M^\downarrow$, we call it a *crash-aware specification*.

Toward defining our linearizability criteria, we start by defining a rewrite system that models the preservation of happens-before ordering from the usual linearizability definition in a more localized way. This formulation has been used in many developments on linearizability [2, 14, 18, 31].

Definition 2.1. We define a string rewrite system \rightsquigarrow with local rewrite rule:

$$s \cdot \alpha:m \cdot \alpha':m' \cdot t \rightsquigarrow s \cdot \alpha':m' \cdot \alpha:m \cdot t$$

whenever $\alpha \neq \alpha'$ and one of the following two conditions hold:

- m and m' are both invocations or both responses, or
- m is an invocation and m' is a response.

The definition of linearizability from the compositional linearizability paper is then given by:

Definition 2.2. A crash-less trace $s \in \mathbb{P}_M^{\text{conc}}$ is linearizable to a crash-less trace $t \in \mathbb{P}_M^{\text{conc}}$ when there exists a sequence of responses $s_P \in M^*$ and a sequence of invocations $s_O \in M^*$ such that $s \cdot s_P \rightsquigarrow t \cdot s_O$. We write $s \rightsquigarrow t$ when s is linearizable to t . We say a crash-less specification ν' linearizes to another one ν , written $\nu' \rightsquigarrow \nu$, when every trace $s \in \nu'$ linearizes to some trace $t \in \nu$.

Note that t is not required to be atomic, as in Herlihy-Wing linearizability, and that s_O is not required to contain every pending invocation of $s \cdot s_P$, unlike most definitions of linearizability. If t is an atomic trace, then this definition is equivalent to the original Herlihy-Wing definition [18, 31].

2.2 Linearizability Under Full-System Crashes

We now define crash-aware linearizability, the criterion we propose in this paper. It requires that each epoch of a trace s linearizes, in the crash-less sense, to the corresponding epoch of t .

Definition 2.3. A crash-aware trace $s \in \mathbb{P}_M^\downarrow$ is *crash-aware linearizable* to a trace $t \in \mathbb{P}_M^\downarrow$ when

$$\|s\| = \|t\| \quad \text{and} \quad \forall i \leq \|s\|. \text{epo}_i(s) \rightsquigarrow \text{epo}_i(t)$$

We denote this as $s \stackrel{\dot{\sim}}{\sim} t$, extending the notation to specifications as with linearizability (Def. 2.2).

Observe that crash-aware linearizability relates crash-aware specifications to crash-aware specifications. This is unusual in the literature on linearizability under crashes, as the other criteria relate a crash-aware specification to a crash-less specification. We discuss the reasons for this later when we have defined two other linearizability criteria and can better compare them.

We now define strict linearizability [2]. Our definition differs from the original one in that it specializes it to full-system crashes (instead of allowing for each agent to crash independently), removes the notion of aborted executions, and generalizes away from atomicity to allow for non-atomic linearized specifications. The first two changes were already considered in Ben-David et al. [3] and make the criterion appropriate for the settings we are interested in, such as NVM and file systems. The later change goes along the lines of the way that Castañeda et al. [7] and Oliveira Vale et al. [31] generalize Herlihy-Wing linearizability [20]. If we restrict our definition so that the linearized trace must be atomic, we obtain the same criterion considered by Ben-David et al. [3].

Definition 2.4. For a crash-aware trace s , we define, whenever well-formed, the crash-less trace

$$\text{ops}(s) := \text{epo}_1(s) \cdot \text{epo}_2(s) \cdot \dots \cdot \text{epo}_{\|s\|}(s)$$

We say a crash-aware trace $s \in \mathbb{P}_M^{\dot{\sim}}$ is *strictly linearizable* to a crash-less trace t , written $s \stackrel{\text{str}}{\sim} t$, when there exists a crash-aware trace t' such that $s \stackrel{\dot{\sim}}{\sim} t'$ and $\text{ops}(t') = t$.

Note that our definition of strict linearizability shows a clear factoring of strict linearizability as crash-aware linearizability followed by crash-removal.

The third and final linearizability criterion we consider here is *durable linearizability* [22]. Durable linearizability is more expressive than strict linearizability [3, 19] in that it considers more objects to be linearizable. This comes at the cost of the extra assumption on the model that new agent names are used in each epoch, which we call the *durability assumption*.

Definition 2.5. We say a crash-aware trace $s \in \mathbb{P}_M^{\dot{\sim}}$ is *durable* when:

$$\forall i, j \leq \|s\|. i \neq j \implies \Upsilon(\text{epo}_i(s)) \cap \Upsilon(\text{epo}_j(s)) = \emptyset$$

where $\Upsilon(t)$ is the set of agents appearing in a trace t . We denote by $\mathbb{P}_M^{\text{dur}} \subseteq \mathbb{P}_M^{\dot{\sim}}$ the subset of well-formed crash-aware traces that are durable.

When a trace s is durable, $\text{ops}(s)$ is always a well-formed crash-less trace. Durable linearizability is then defined in terms of usual crash-less linearizability. Our definition, similarly to our definitions of the other linearizability criteria we consider, generalizes away from atomicity by allowing linearized traces to be non-atomic and allows for the specification of blocking objects, as it does not require all uncompleted pending invocations to be removed. It is, however, fully equivalent to the original definition of durable linearizability if we require that the linearized trace be atomic.

Definition 2.6. We say a durable trace $s \in \mathbb{P}_M^{\text{dur}}$ is *durably linearizable*, written $s \stackrel{\text{dur}}{\sim} t$, to a crash-less trace t when $\text{ops}(s) \rightsquigarrow t$.

Note that durable linearizability corresponds to the inverse factoring to strict linearizability, one first removes crashes and then uses crash-less linearizability. These two factorings play an important technical role in our proofs. Moreover, it is possible to show that both criteria factor (in a different sense) through crash-aware linearizability.

PROPOSITION 2.7.

- If $s' \stackrel{\dot{\sim}}{\sim} s$ and $s \stackrel{\text{str}}{\sim} t$ then $s' \stackrel{\text{str}}{\sim} t$
- If $s' \stackrel{\dot{\sim}}{\sim} s$ and $s \stackrel{\text{dur}}{\sim} t$ then $s' \stackrel{\text{dur}}{\sim} t$

Because of this fact, in practice, when verifying durably linearizable objects, we find it useful to use a crash-aware specification v^{mid} satisfying: $v' \xrightarrow{\text{c}} v^{\text{mid}}$ and $v^{\text{mid}} \xrightarrow{\text{d}} v$. This allows us to consider less concurrent traces within the linearized specification for v' by linearizing as much as possible within each epoch of v^{mid} first. This allows us to obtain the benefits of both crash-aware and durable linearizability simultaneously: by maintaining both v^{mid} and v we can still express durably linearizable specifications, but by manipulating v^{mid} we achieve the same level of compositionality as crash-aware linearizability. This technique is not necessary for strict linearizability because we can just use crash-aware linearizability directly by always picking v^{mid} so that $\text{ops}(v^{\text{mid}}) = v$.

2.3 Specifying a Buffered Memory Cell

In §1 we mentioned that we use crash-aware linearizability to specify a buffered memory cell with signature BCell. As an example, we define here what the linearized specification for a buffered memory cell implementation would be under crash-aware linearizability.

An example of a trace of a concrete buffered memory cell v'_{BCell} is:

$$\alpha_1:\text{store}(1) \cdot \alpha_2:\text{load}() \cdot \alpha_1:\text{ok} \cdot \alpha_2:0 \cdot \alpha_2:\text{flush}() \cdot \alpha_1:\text{store}(2) \cdot \alpha_1:\text{ok} \cdot \frac{1}{2} \cdot \alpha_3:\text{load}() \cdot \alpha_3:1$$

The trace above is crash-aware linearizable to the following trace, among others:

$$\alpha_2:\text{load}() \cdot \alpha_2:0 \cdot \alpha_1:\text{store}(1) \cdot \alpha_1:\text{ok} \cdot \alpha_2:\text{flush}() \cdot \alpha_2:\text{ok} \cdot \alpha_1:\text{store}(2) \cdot \alpha_1:\text{ok} \cdot \frac{1}{2} \cdot \alpha_3:\text{load}() \cdot \alpha_3:1$$

We specify the semantics of the buffered memory cell by a set of traces v'_{BCell} with only events that are allowed by the signature BCell. Crashes can happen at any point. To specify the correctness of v'_{BCell} we require it to be crash-aware linearizable to the atomic linearized specification v_{BCell} . Because we show observational refinement, we are able to leave v'_{BCell} unspecified for the sake of verifying the FLiT implementation, as only the linearized specification will be necessary. The linearized specification v_{BCell} is then defined by:

$$s \in v_{\text{BCell}} \iff s \text{ is atomic } \wedge (\forall s_1, s_2. \forall v. s = s_1 \cdot \alpha:\text{load}() \cdot \alpha:v \cdot s_2 \implies v \in \text{snd}(\text{mstate}(s_1)))$$

where $\text{mstate}(s)$ assigns to an atomic complete trace s a set of pairs $\text{mstate}(s) \subseteq \text{Val} \times \text{Val}$. A pair $(v_p, v_b) \in \text{mstate}(s)$ consists of a possibility for a value v_p that has persisted and a value v_b currently in the buffer. $\text{mstate}(s)$ is then the function inductively defined below ($v_0 \in \text{Val}$ is an identified initial value for the memory cell):

$$\text{mstate}(\epsilon) := \{(v_0, v_0)\} \quad \text{mstate}(s \cdot \frac{1}{2}) := \{(v, v) \mid \exists v'. (v, v') \in \text{mstate}(s)\}$$

$$\text{mstate}(s \cdot \alpha:\text{store}(v) \cdot \alpha:\text{ok}) := \{(v', v) \mid \exists v''. (v', v'') \in \text{mstate}(s)\} \cup \{(v, v)\}$$

$$\text{mstate}(s \cdot \alpha:\text{load}() \cdot \alpha:v) := \{(v', v) \mid (v', v) \in \text{mstate}(s)\}$$

$$\text{mstate}(s \cdot \alpha:\text{flush}() \cdot \alpha:\text{ok}) := \{(v, v) \mid (v', v) \in \text{mstate}(s)\}$$

The function $\text{snd}(p)$ projects into the second component v_b of the pair $p = (v_p, v_b)$, so that $\text{snd}(\text{mstate}(s)) = \{v_b \mid \exists v_p. (v_p, v_b) \in \text{mstate}(s)\}$.

Note that we could have specified it instead using a labeled state transition system (LTS), in which case v_{BCell} is the set of traces that start from the initial state of the LTS. Either way of defining v_{BCell} defines the same set of traces.

2.4 Contrasting Crash-Aware Linearizability

We now compare crash-aware linearizability against strict and durable linearizability. We will not compare strict and durable linearizability against each other since they are not new to our work, and refer the interested reader to the following references [3, 19, 22]. We do briefly mention a key difference that applies to crash-aware linearizability as well. In strict and crash-aware linearizability, a pending invocation must be linearized *within* the epoch it was issued. Durable linearizability,

however, allows for a pending invocation to be linearized in (essentially) a later epoch by allowing those pending invocations to be reordered after events from later epochs. This is what makes it more expressive than strict linearizability, allowing for more complex crash behaviors, such as recovering parts of a data structure only when they are demanded by a client, which could happen several epochs later. As explained in the remark at the end of §2.2 crash-aware linearizability interacts well with durable linearizability. This will be discussed further in §5.3.

As we saw, both durable and strict linearizability factor through crash-aware linearizability. The key difference between the two former criteria and the latter is that the former use crash-less linearized specifications, while the latter uses crash-aware linearized specifications. So let's refer to the former as *crash-unaware* criteria.

Crash-unaware criteria reduce the correctness of an object with crashes to that of an object without crashes. This makes them great at specifying objects with very strong persistency guarantees, that is, objects whose whole state (or almost) persists after a crash. But it makes them quite deficient at specifying objects with weaker persistency guarantees such as volatile objects (all of the state is lost on a crash), objects with hybrid persistency (part of the state is volatile and part of the state is persistent), or objects whose persistency features some degree of non-determinism (such as in buffered memory). Some of these issues were already known. For instance, in the original durable linearizability paper [22], it is noted that the criteria do not behave well when used to specify a buffered object, requiring them to define an *ad-hoc* notion of *buffered* durable linearizability which does not satisfy locality, making it not compositional.

Consider the simple problem of specifying the correctness of a volatile object. Given a crash-less specification v , we can construct a crash-aware specification $\text{vol}(v)$ of a volatile version of that object by the Kleene algebra formula $\text{vol}(v) := (v \cdot \frac{1}{2})^* \cdot v$.

For example, given the usual atomic counter specification v_{Counter} , the following trace is allowed by $\text{vol}(v_{\text{Counter}})$ (the subscript 1 under the Counter operations will be useful later):

$$s_1 = \alpha_1:\text{inc}_1 \cdot \alpha_1:\text{ok} \cdot \alpha_2:\text{get}_1 \cdot \alpha_2:1 \cdot \frac{1}{2} \cdot \alpha_3:\text{get}_1 \cdot \alpha_3:0$$

Note that the crash move $\frac{1}{2}$ plays a crucial role in the specification, as the counter only resets to 0 after a crash event (such as the last get event in s_1), making the linearized specification deterministic.

Under crash-aware linearizability, a concurrent object v' correctly implements a volatile version of a crash-less object v when $v' \xrightarrow{\frac{1}{2}} \text{vol}(v)$. With our methods, it is easy to show that

PROPOSITION 2.8. *If $v' \rightsquigarrow v$ then $\text{vol}(v') \xrightarrow{\frac{1}{2}} \text{vol}(v)$.*

A consequence of this is that if we have an implementation M that implements a crash-less object linearizable to v_B on top of a crash-less object v_A , then if we run M on each epoch on top of $\text{vol}(v_A)$ then M implements an object crash-aware linearizable to $\text{vol}(v_B)$. Formally,

$$v_A; M \rightsquigarrow v_B \implies \text{vol}(v_A); \text{vol}(M) \xrightarrow{\frac{1}{2}} \text{vol}(v_B)$$

In other words, crash-aware linearizability is able to appropriately specify and characterize the correctness of volatile objects in a way that is useful to a client.

Now, consider what happens if we try to specify a volatile object using a crash-unaware criterion. A crash-unaware linearized specification will need to include both of the following traces:

$$\text{ops}(s_1) = \alpha_1:\text{inc}_1 \cdot \alpha_1:\text{ok} \cdot \alpha_2:\text{get}_1 \cdot \alpha_2:1 \cdot \alpha_3:\text{get}_1 \cdot \alpha_3:0 \quad \text{and} \quad \alpha_1:\text{inc}_1 \cdot \alpha_1:\text{ok} \cdot \alpha_2:\text{get}_1 \cdot \alpha_2:1 \cdot \alpha_3:\text{get}_1 \cdot \alpha_3:1$$

so that the linearized specification under crash-unaware criteria must admit non-deterministically resetting the counter at any point. This can happen at any point, but the point at which it happens is not detectable in the linearized specification, which makes the specification quite weak. This means that even if some observational refinement theorem (*à la* Filipovic et al. [14]) holds for the crash-unaware criterion, the client to the linearized specification will need to contend with non-determinism, making the contextual refinement, and hence vertical composition, weaker.

This issue is compounded when considering horizontal composition. Both durable and strict linearizability are known to satisfy locality. However, those locality theorems introduce even more non-determinism into the resulting linearized specifications. Consider now a second trace s_2 for a second counter independent of the counter in play s_1

$$s_2 = \alpha_1:\text{inc}_2 \cdot \alpha_1:\text{ok} \cdot \alpha_2:\text{get}_2 \cdot \alpha_2:1 \cdot \zeta \cdot \alpha_3:\text{get}_2 \cdot \alpha_3:0$$

Any trace in their parallel composition $s_1 \otimes s_2$ (the set of well-formed crash-aware interleavings of s_1 and s_2 , defined in §3) synchronizes on the crash, so both counters reset their state at the same time. For example, the following crash-aware trace belongs to $s_1 \otimes s_2$:

$$\alpha_1:\text{inc}_1 \cdot \alpha_1:\text{ok} \cdot \alpha_1:\text{inc}_2 \cdot \alpha_1:\text{ok} \cdot \alpha_2:\text{get}_1 \cdot \alpha_2:1 \cdot \alpha_2:\text{get}_2 \cdot \alpha_2:1 \cdot \zeta \cdot \alpha_3:\text{get}_2 \cdot \alpha_3:0 \cdot \alpha_3:\text{get}_1 \cdot \alpha_3:0$$

Meanwhile, the corresponding linearized specifications under durable or strict linearizability include these traces without the crash event, i.e., $\text{ops}(s_1)$ and $\text{ops}(s_2)$. Hence, the following trace is in their parallel composition $\text{ops}(s_1) \otimes \text{ops}(s_2)$ (the set of their well-formed crash-less interleavings):

$$\alpha_1:\text{inc}_1 \cdot \alpha_1:\text{ok} \cdot \alpha_1:\text{inc}_2 \cdot \alpha_1:\text{ok} \cdot \alpha_2:\text{get}_1 \cdot \alpha_2:1 \cdot \alpha_3:\text{get}_1 \cdot \alpha_3:0 \cdot \alpha_2:\text{get}_2 \cdot \alpha_2:1 \cdot \alpha_3:\text{get}_2 \cdot \alpha_3:0$$

There is no trace of the concrete horizontally composed volatile counters that is linearizable to the trace above, as we must at least introduce a crash right before $\alpha_3:\text{get}_1$ to justify its return $\alpha_3:0$:

$$\alpha_1:\text{inc}_1 \cdot \alpha_1:\text{ok} \cdot \alpha_1:\text{inc}_2 \cdot \alpha_1:\text{ok} \cdot \alpha_2:\text{get}_1 \cdot \alpha_2:1 \cdot \zeta \cdot \alpha_3:\text{get}_1 \cdot \alpha_3:0 \cdot \alpha_2:\text{get}_2 \cdot \alpha_2:1 \cdot \alpha_3:\text{get}_2 \cdot \alpha_3:0$$

This makes the trace inconsistent with the semantics of the second counter, as the crash should also have reset it, so that $\alpha_2:\text{get}_2$ should not return $\alpha_2:1$. The same kind of argument shows how crash-unaware criteria fail to accurately handle hybrid and buffered objects (all the traces above are valid for a buffered counter, for example).

3 A CONCURRENT GAME SEMANTICS WITH CRASHES

So far, in §2 we focused on three linearizability criteria in a unstructured setup. For instance, we did not enforce typing on specifications. This will not be enough to achieve the degree of compositionality we seek, especially as we treat objects as open components.

In this section, we discuss our compositional model with crashes in detail. The model is defined using a simple game semantics. The reader not familiar with game semantics jargon will find the following approximation useful. A *game* A, B roughly corresponds to a type; a *move* of the game A corresponds to an event of type A which also has a *polarity*, i.e. its metadata (such as the name of the agent who issued it, and whether it is a move by the environment or by the system); a *play* over a game A is a trace of that type. Crucially, plays can have higher-order types (unlike in most trace semantics); in particular, we may form the affine implication game $A \multimap B$ (the type of code using an object of type A to implement one of type B) whose plays are well-formed traces involving moves from both A and B ; a *strategy* σ of type A is the denotation of some computation, be it a state transition system, or the semantics of some code. It is represented as some prefix-closed set of plays² of its type A . Readers looking for comprehensive introductions to game semantics may benefit from Abramsky and McCusker [1], Ghica [15], Hyland [21] though we warn that our model simplifies several aspects of these game semantics, which are not necessary for our purposes.

3.1 Games with Full-System Crashes

Definition 3.1 (Polarities and Moves). A *move set* consists of a set of *moves* M together with an assignment $\lambda : M \rightarrow \sum_{\alpha \in \Gamma} \{O, P\}$, that is, every move is labeled with the agent who plays it and

²It is folklore that prefix-closed sets of traces are in one-to-one correspondence with equivalence classes of transition systems under forward-backward simulation [27]. Therefore, all of our results translate to equivalent statements that hold up to forward-backward simulation. We use a presentation based on prefix-closed sets of traces as it aligns well with the typical treatment of linearizability, while simplifying many aspects of the presentation and of the compositional structure.

whether or not it is an environment (O) or a system (P) move. The elements of $\sum_{\alpha \in \Upsilon} \{O, P\}$ are called polarities and are denoted by $\alpha:O$ or $\alpha:P$.

Most of the games we use in practice will be defined by first providing an effect signature. An effect signature is a collection of operations, or effects, $E = (e_i)_{i \in I}$ together with assignments $\text{par}(-), \text{ar}(-) : E \rightarrow \text{Set}$ of a set of parameters $\text{par}(e)$ and a set of return values $\text{ar}(e)$ for each operation $e \in E$. This is conveniently described by the following notation.

$$E = \{e_i : \text{par}(e_i) \rightarrow \text{ar}(e_i) \mid i \in I\}$$

All the signatures defined in §1 are effect signatures. We call an Υ -indexed collection of effect signatures $E = (E[\alpha])_{\alpha \in \Upsilon}$ a concurrent effect signature. Given a concurrent effect signature E we define a corresponding move set as follows:

$$M_{\dagger E} := \sum_{\alpha \in \Upsilon} (\sum_{e \in E[\alpha]} \text{par}(e) + \sum_{e \in E[\alpha]} \text{ar}(e))$$

$\lambda_{\dagger E}(\alpha:e(a)) := \alpha:O, e \in E[\alpha] \wedge a \in \text{par}(e)$ $\lambda_{\dagger E}(\alpha:v) := \alpha:P, v \in \text{ar}(e)$ for some $e \in E[\alpha]$
in other words, moves in $M_{\dagger E}$ are either $\alpha:e(a)$ for $e \in E[\alpha]$ and $a \in \text{par}(e)$, in which case $\lambda_{\dagger E}(\alpha:e(a)) = \alpha:O$, or $\alpha:v$ with $v \in \text{ar}(e)$, in which case $\lambda_{\dagger E}(\alpha:v) = \alpha:P$.

Definition 3.2. We denote a crash by \downarrow . Given a move set M we write M^{\downarrow} for its extension $M + \{\downarrow\}$ with a crash move. We also extend its polarity function λ into λ^{\downarrow} with the assignment $\lambda^{\downarrow}(\downarrow) = \downarrow$.

Recall that given a sequence $s \in M^*$, we write $\pi_{\alpha}(s)$ for the projection of s to its largest subsequence involving only events by $\alpha \in \Upsilon$.

Definition 3.3. A game $A = (M_A, \lambda_A, P_A)$ consists of a move set (M_A, λ_A) and a non-empty, prefix-closed set of well-formed crash-less plays $P_A \subseteq \mathbb{P}_{M_A}^{\text{conc}}$ satisfying $P_A = \|\alpha \in \Upsilon \pi_{\alpha}(P_A)$. We write $P_A^{\downarrow} \subseteq \mathbb{P}_{M_A}^{\downarrow}$ for the set $P_A^{\downarrow} := (P_A \cdot \downarrow)^* \cdot P_A$.

The set of plays P_A of a game A defines which plays are valid plays within an epoch. It is required to be an arbitrary parallel compositions of the sequential plays that each agent can perform. Meanwhile, P_A^{\downarrow} , the corresponding set of crash-aware plays, is defined by simply allowing crashes to happen at any point in an epoch.

Some examples of crash-aware games are now due. The simplest game is the game $\mathbf{1} := (\emptyset, \emptyset, \{\epsilon\})$. The game $\mathbf{1}$ has no non-crash moves, and its only crash-aware plays are the empty sequence ϵ and sequences of crashes $\downarrow \cdot \downarrow \cdot \dots \cdot \downarrow$.

Another game is the game $\Sigma = ((\sum_{\alpha \in \Upsilon} q + a), (\sum_{\alpha \in \Upsilon} \alpha:O + \alpha:P), \|\alpha \in \Upsilon \downarrow\{q \cdot a\})$ where $\downarrow -$ stands for prefix-closure. Unrolling this definition, every agent has two moves: an O -move q (question) and a P -move a (answer). The only valid sequential plays are $q \cdot a$ and its prefixes, and the valid plays for the game are interleavings of these sequential plays at each epoch, such as:

$$\alpha:q \cdot \alpha':q \cdot \alpha:a \cdot \downarrow \cdot \alpha':q \cdot \alpha:q \cdot \alpha:a \cdot \downarrow \cdot \alpha':q$$

The most important kind of game for our examples are games $\dagger E$ generated by effect signatures E . We can extend the move set $(M_{\dagger E}, \lambda_{\dagger E})$ into a game with set of valid plays $P_{\dagger E}$ defined by:

$$P_{\dagger E} := \|\alpha \in \Upsilon \downarrow (\cup_{e \in E[\alpha]} \cup_{a \in \text{par}(e)} \cup_{v \in \text{ar}(e)} \alpha:e(a) \cdot \alpha:v)^*$$

That is to say, locally, each agent $\alpha \in \Upsilon$ is allowed to alternate between making a call to an effect $e(a)$ in $E[\alpha]$ or providing a response to the previously issued effect. For instance, recall that we defined, in §1, a signature `BCell` encoding the operations available to a buffered memory cell. This defines a concurrent effect signature $\text{BCell}[\alpha] = \text{BCell}$. The corresponding set of valid crash-aware plays $P_{\dagger \text{BCell}}$ includes all traces seen in §2.3.

3.2 Combining Games

We now define a few combinators on games. We start by defining a dualizing operation on move sets, which swaps the role of environment and system.

Definition 3.4 (Dual Move Set). Given a move set (M, λ) we define the moveset (M^\perp, λ^\perp) by $M^\perp := M$ and $\lambda^\perp(m) := \lambda(m)^\perp$, where $(\alpha:O)^\perp := \alpha:P$ and $(\alpha:P)^\perp := \alpha:O$.

In the context of games A, B, C , given $s \in \mathbb{P}_{M_A+M_B}^\perp$ we define $s \upharpoonright_{A,-} \in \mathbb{P}_{M_A}^\perp$ and $s \upharpoonright_{-,B} \in \mathbb{P}_{M_B}^\perp$ to be the projections to the corresponding components of $M_A + M_B$, but keeping the crash moves in the projections too. Similarly, given $s \in \mathbb{P}_{M_A+M_B+M_C}^\perp$, we write $s \upharpoonright_{A,B,-}$, for the projection of s to its largest subsequence with only moves in A, B and crashes; we similarly define $s \upharpoonright_{A,-,C}$ and $s \upharpoonright_{-,B,C}$. We now define horizontal composition of games, and the affine arrow.

Definition 3.5. Fix games A and B . We define the games $A \otimes B$ and $A \multimap B$ by the following data

$$\begin{aligned} M_{A \otimes B} &:= M_A + M_B & \lambda_{A \otimes B} &:= \lambda_A + \lambda_B & P_{A \otimes B} &:= \{s \in \mathbb{P}_{M_A+M_B} \mid s \upharpoonright_{A,-} \in P_A \wedge s \upharpoonright_{-,B} \in P_B\} \\ M_{A \multimap B} &:= M_A^\perp + M_B & \lambda_{A \multimap B} &:= \lambda_A^\perp + \lambda_B & P_{A \multimap B} &:= \{s \in \mathbb{P}_{M_A^\perp+M_B} \mid s \upharpoonright_{A,-} \in P_A \wedge s \upharpoonright_{-,B} \in P_B\} \end{aligned}$$

It is implicit in this definition that by composing in parallel the two crash-aware plays, the resulting set of traces synchronizes the crash events, merging them into a single crash event and then producing any (locally sequential) parallel composition of the subtraces appearing in each epoch. Consider, for instance, the two plays below on the left, each of type Σ :

$$\begin{array}{cccccc} \alpha:q & \alpha:a & \downarrow & \alpha:q & \alpha:a & \downarrow & \alpha':q & \alpha':a & & \alpha:q & \alpha:a & | & \alpha:q & \alpha:a & | & \alpha':q & \alpha':a \\ & & & \otimes & & & & & & \otimes & & \downarrow & \otimes & & \downarrow & \otimes & \\ \alpha:q & \alpha:a & \downarrow & \alpha':q & \alpha':a & \downarrow & \alpha':q & \alpha':a & & \alpha:q & \alpha:a & | & \alpha':q & \alpha':a & | & \alpha':q & \alpha':a \end{array} \implies$$

The resulting set of traces synchronizes the crashes, as depicted on the right. For example, the following is a valid trace in their horizontal composition:

$$\alpha:q \cdot \alpha:a \cdot \alpha:q \cdot \alpha:a \cdot \downarrow \cdot \alpha:q \cdot \alpha':q \cdot \alpha:a \cdot \alpha':a \cdot \downarrow \cdot \alpha':q \cdot \alpha':a \cdot \alpha':q \cdot \alpha':a$$

Similarly, consider the following play s of $\Sigma \multimap \Sigma$ (on the left):

$$\begin{array}{ccc|c} \Sigma & \alpha:q & \alpha':q & \alpha':q \\ \uparrow & & & \\ \Sigma & \alpha:q & \alpha:a & \alpha':q \end{array} \quad \left| \quad \begin{array}{l} s \upharpoonright_{-, \Sigma} = \alpha:q \cdot \alpha':q \cdot \downarrow \cdot \alpha':q \\ s \upharpoonright_{\Sigma, -} = \alpha:q \cdot \alpha:a \cdot \downarrow \cdot \alpha':q \end{array} \right.$$

or, depicted sequentially: $\alpha:q \cdot \alpha:q \cdot \alpha':q \cdot \alpha:a \cdot \downarrow \cdot \alpha':q \cdot \alpha':q$. Note that the crash signal synchronizes across the source and target components of the play. This models that the crashes are *synchronous* across components (they happen in all components at once) and that they are *instantaneous* (it takes negligible time for the crash to propagate to components). On the right, above, we see the projections of s to the source and target components. Importantly, the crash event is retained in both projections, so it effectively belongs to both components.

3.3 Crash-Aware Strategies

We now define strategies, which are the denotations of both object specifications and code.

Definition 3.6 (Crash-Aware Strategy). A crash-aware strategy $\sigma : A$ over a game A is a non-empty, prefix-closed, subset $\sigma \subseteq P_A^\perp$, which is moreover \downarrow -receptive in that

$$\forall s \in \sigma. s \cdot \downarrow \in P_A^\perp \implies s \cdot \downarrow \in \sigma$$

\downarrow -receptivity models the usual assumption that crashes may non-deterministically happen at any point in an execution. It plays a crucial role in proving the locality property.

We specify the semantics of objects using strategies. For example, in §2.3 we specified the linearized buffered memory cell by a strategy $\nu_{\text{BCell}} : \dagger\text{BCell}$. The denotations of implementations, such as $M_{\text{FLIT}} : \dagger\text{BCell} \otimes \dagger\text{Counter} \multimap \dagger\text{FLIT}$ or $M_{\text{Snapshot}} : \dagger\text{Mem} \multimap \dagger\text{Snapshot}$ mentioned in §1 are examples of strategies with the affine arrow type. Strategies of type $A \multimap B$ can be vertically composed, which amounts to the usual motto of “interaction + hiding”.

Definition 3.7. Given strategies $\sigma : A \multimap B$ and $\tau : B \multimap C$ we define their vertical composition $\sigma; \tau : A \multimap C$ by: $\sigma; \tau := \{s \upharpoonright_{A,-,C} \in \mathbb{P}_{A \multimap C}^{\dagger} \mid \exists s \in ((M_A + M_B + M_C)^{\dagger})^* . s \upharpoonright_{A,B,-} \in \sigma \wedge s \upharpoonright_{-,B,C} \in \tau\}$

PROPOSITION 3.8. *Composition of crash-aware strategies is well-defined and associative.*

For the reader with familiarity with category theory, we can package all the information above:

Definition 3.9. We denote by **Crash** the semicategory of crash-aware games, with crash-aware strategies $\sigma : A \multimap B$ as morphisms between games A and B , and composition given by $;$.

Unfortunately, and this is a common phenomenon in concurrency models, **Crash** does not assemble into a category, as the vertical composition operation $;$ does not have a neutral element. That is, to say, there is no choice of strategies $\text{id}_A : A \multimap A$ for which $\text{id}_A; \sigma; \text{id}_B = \sigma$ for every $\sigma : A \multimap B$. This issue is explained extensively in Oliveira Vale et al. [31] in the context of concurrent games.

We follow the approach from compositional linearizability, and start by noting that there are obvious candidates $\text{crashcopy}_A : A \multimap A$ for the neutral elements, which are called the copycat strategies and formalize the code seen in Fig. 3. The copycat strategy is *idempotent*, in that for all games A , $\text{crashcopy}_A; \text{crashcopy}_A = \text{crashcopy}_A$. This essentially means that the crashcopy_A at least behaves like a neutral element for itself. In fact, they behave like a neutral element with respect to any strategy which is a parallel composition of sequential strategies. This fact justifies defining a class of strategies that behaves well when composed with the copycat:

Definition 3.10. We say a strategy $\sigma : A \multimap B$ is saturated with respect to crashcopy when

$$\text{crashcopy}_A; \sigma; \text{crashcopy}_B = \sigma$$

Since crashcopy is idempotent, it is saturated. Moreover, by definition, crashcopy behaves as a neutral element for strategy composition of saturated strategies. It is also easy to see that saturated strategies compose. Note that this means that we can promote **Crash** to a category by restricting attention to these saturated strategies.

Saturation for concurrent strategies corresponds to, beyond O -receptivity (the environment can make valid moves whenever it wants), strategies that are insensitive to certain delays, which might be caused, for instance, if an agent is preempted. This is typically formalized using the rewrite system \rightsquigarrow we defined in §2.1 and redefined now in light of our more detailed formalism.

Definition 3.11. Let $A = (M_A, P_A)$ be a game. We define a string rewrite system $(P_A, \rightsquigarrow_A)$ with local rewrite rules:

- $\forall m, m' \in M_A. \forall X \in \{O, P\}. \lambda_A(m) = \alpha : X \wedge \lambda_A(m') = \alpha' : X \wedge \alpha \neq \alpha' \implies m \cdot m' \rightsquigarrow_A m' \cdot m$
- $\forall m, m' \in M_A. \lambda_A(m) = \alpha : O \wedge \lambda_A(m') = \alpha' : P \wedge \alpha \neq \alpha' \implies m \cdot m' \rightsquigarrow_A m' \cdot m$

Fig. 3. Code corresponding to the copycat strategy $\text{crashcopy}_{\dagger F \multimap \dagger F} : \dagger F \multimap \dagger F$

A concrete characterization of saturation for crash-aware strategies is possible, but we do not cover it here for the sake of space (see the appendix). We will soon see an equivalent characterization in terms of crash-aware linearizability, which will be sufficient for our purposes.

3.4 Refinement and Horizontal Composition

Before proceeding, we briefly address refinement and horizontal composition. We take as our notion of refinement behavior containment, $\sigma \subseteq \tau$, with joins given by set union. This makes all of the models we discuss into enriched (semi)categories over join semi-lattices, which means that:

PROPOSITION 3.12. *Strategy composition $-; -$ is monotonic and join-preserving.*

For horizontal composition, recall that we have already defined a game $A \otimes B \in \underline{\text{Crash}}$. The tensor can be extended to strategies $\sigma : A \multimap B$ and $\tau : A' \multimap B'$ by:

$$\sigma \otimes \tau := \{s \in P_{A \otimes A' \multimap B \otimes B'} \mid s \upharpoonright_{A \multimap B} \in \sigma \wedge s \upharpoonright_{A' \multimap B'} \in \tau\}$$

PROPOSITION 3.13. *Let Crash be the restriction of the semicategory $\underline{\text{Crash}}$ to strategies saturated with respect to crashcopy. Then, $(\text{Crash}, - \otimes -, 1)$ defines an enriched symmetric monoidal category.*

These definitions permit us to prove Prop. 3.13. This means that $- \otimes -$ defines a monotonic and join-preserving functor so that horizontal composition behaves well with respect to both vertical composition and refinement. This formalizes what we mean when we say that our model is compositional. It remains to extend this compositional structure to linearizability.

4 THREE LINEARIZABILITY CRITERIA REVISITED

We now revisit the linearizability criteria discussed in §2 from the perspective of our just defined model and following ideas from compositional linearizability. In particular, we argue that their methodology recovers crash-aware linearizability as the notion of linearizability associated with the compositional structure of our model and use their general theorem around locality and observational refinement to obtain these results for crash-aware linearizability. Then, we extend these results to strict and durable linearizability by analyzing translations from our crash-aware model to the crash-less model from compositional linearizability.

4.1 Abstract Crash-Aware Linearizability

In §2 we defined a new linearizability criterion which we called *crash-aware linearizability* ($\overset{\zeta}{\sim}$). We, however, did not come up with this definition of linearizability. Instead, following the methodology of compositional linearizability, we have derived it from the structure of the model, $\underline{\text{Crash}}$.

To understand this, we start by defining the operation $K_\zeta - : \underline{\text{Crash}} \rightarrow \text{Crash}$ by the formula

$$K_\zeta \tau := \text{crashcopy}_A; \tau; \text{crashcopy}_B$$

for $\tau : A \multimap B \in \underline{\text{Crash}}$. This operation assigns to τ the smallest saturated strategy containing τ .

The framework of compositional linearizability proposes that the native notion of linearizability for crash-aware objects should be equivalent to the refinement $v' \subseteq K_\zeta v$. Indeed, we are able to show the following characterization of $K_\zeta -$, which provides a concrete characterization of $K_\zeta v$ as the set of all plays that are crash-aware linearizable w.r.t. v .

PROPOSITION 4.1. *For any crash-aware strategy $v : A$ it holds that: $K_\zeta v = \{s \in \mathbb{P}_A^\zeta \mid \exists t \in v. s \overset{\zeta}{\sim} t\}$.*

It follows immediately from this characterization that

PROPOSITION 4.2. *v' is crash-aware linearizable w.r.t. v if and only if $v' \subseteq K_\zeta v$.*

This effectively turns linearizability into a refinement property. This has many benefits from the point of view of verification, as refinement techniques are well-understood. Moreover, since we derive it in this way, we may use the general category-theoretic result in Oliveira Vale et al. [31] to obtain locality and observational refinement.

PROPOSITION 4.3 (OBSERVATIONAL REFINEMENT AND LOCALITY).

- $v'_A : A$ is crash-aware linearizable w.r.t $v_A : A$ iff for all saturated $\sigma : A \multimap B$, $v'_A; \sigma \subseteq v_A; \sigma$
- For $v'_A : A, v'_B : B$ and $v_A : A, v_B : B$: $v'_A \xrightarrow{\xi} v_A$ and $v'_B \xrightarrow{\xi} v_B$ if and only if $v'_A \otimes v'_B \xrightarrow{\xi} v_A \otimes v_B$

4.2 Compositional Verification of a File System Fragment

To showcase the benefits of compositionality and to show that crash-aware linearizability provides a flexible criterion for mixing objects with different, and complicated, crash behaviors, we verify against a crash-aware linearizable specification a fragment of a file API. Instead of providing a detailed description, we emphasize the salient aspects to our point (a detailed description is available in the appendix). The system also features recovery, our handling of which is discussed later (§5).

The file system fragment involves four main objects: the file interface File, a disk interface Disk implemented using a disk array Disk[N] with a finite number N of disks each with $S + 1$ blocks, and a write-ahead log Log. The signatures for File and Disk are given below:

$$\text{File} := \left\{ \begin{array}{l} \text{write} : \text{block_id} \times \text{file_id} \times \text{block} \rightarrow 1, \\ \text{read} : \text{block_id} \times \text{file_id} \rightarrow \text{block}, \\ \text{swap} : \text{block_id} \times \text{block_id} \times \text{file_id} \times \text{file_id} \rightarrow 1 \end{array} \right\} \quad \text{Disk} := \left\{ \begin{array}{l} \text{write} : \text{block_id} \times \text{block} \rightarrow 1, \\ \text{read} : \text{block_id} \rightarrow \text{block} \end{array} \right\}$$

The file interface exposes a two-level structure. At the first level lies a set of folders, each occupying a single disk block as its inode. For simplicity, the API uses block ids instead of strings to uniquely identify folders. Each folder contains a set of files identified by their file id. The swap operation swaps the pointers in the respective folders' inodes, which provides a symmetric move operation as seen in actual file systems. The write and read operations are as usual. The file interface is implemented on top of a disk, providing write and read operations to read and write to a block.

All the objects involved are specified using crash-aware linearizability. For instance, a single disk is specified as the horizontal composition of its blocks, using locality, guaranteeing that its concrete specification $v'_{\text{Disk}} : \dagger\text{Disk}$ is crash-aware linearizable to a specification $v_{\text{Disk}} : \dagger\text{Disk}$ which guarantees read and writes are persistent and atomic. The disk array specification $v'_{\text{Disk}[N]}$ is required to be crash-aware linearizable to the horizontal composition of N atomic disk specifications $v_{\text{Disk}[N]} := \otimes_{i \in [N]} v_{\text{Disk}}$. The concrete object v'_{File} is required to be crash-aware linearizable to a specification v_{File} that ensures that writes, reads and swaps are persistent and seem to happen atomically. All the specifications also enforce that the recovery routines correctly reconstruct any relevant lost state after a crash.

We implement the replicated disk on top of the disk array by replicating writes to all the disks in the array in a specific order. Reads to the disk array non-deterministically choose a disk to read from, mimicking the behavior of a disk array controller. On a crash, a recovery procedure copies the contents of the first disks to all disks.

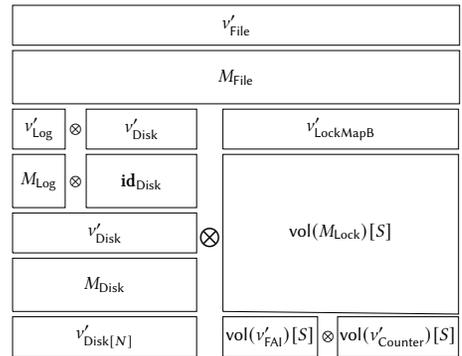


Fig. 4. The structure of our File example.

The File implementation M_{File} for write and read is mostly straight-forward. The swap operation requires special treatment for its atomicity. As swap operations need to update two different folders (and thus two different disk blocks), to ensure persistency, we record the operations in a write-ahead log v'_{Log} so that the recovery routine can finish incomplete operations. The log is itself implemented on top of a single block of the disk together with a volatile array and a volatile lock (omitted from Fig. 4). Since the disk is itself equivalent to the parallel composition of individual blocks, we use locality together with our compositionality properties (the symmetric monoidal structure of the model) to separate the part of the disk used for the log, from the rest of the disk.

The file system also uses a set of dynamically allocated locks v'_{LockMapB} to guarantee atomicity when writing to a block. These locks are volatile objects residing in memory that only last for the duration of a single File operation. Because of this, we use the verified lock from Oliveira Vale et al. [31] and lift it to a volatile object using Prop. 2.8, benefiting from the fact that we have connected our model to their model. The whole structure of the example is depicted in Fig. 4.

At this point it is worth remarking that even this small fragment of a file system features a mix of persistent objects, volatile objects, and objects that fit neither category well. Some of the objects involved are horizontal compositions of these objects, making them have hybrid crash behavior. We model all of these objects using crash-aware linearizability, which proves to be robust enough to verify the whole system compositionally.

4.3 Crash Abstraction

Recall that strict and durable linearizability relate a crash-aware concrete specification to a crash-less specification. In this section we develop conversions between these computational models, which serve as a building block for strict and durable linearizability. So, first, we briefly recall that:

Definition 4.4. Given a game A , a crash-less strategy $\sigma : A$ consists of a non-empty, prefix-closed set of plays $\sigma \subseteq P_A$.

The main difficulty in removing crashes from a play s is that the removal may generate traces that do not satisfy well-formedness. This happens when the same agent has a pending invocation in one epoch and also moves in a later epoch. So, in the definition of the operation $-^b$ (read *de-crash*, and the same as $\text{ops}(-)$), the projections $\text{ops}(s)$ are required to be well-formed plays.

Definition 4.5. Given a game $A = (M_A, \lambda_A, P_A)$ we define the game A^b , by:

$$M_{A^b} := M_A \quad \lambda_{A^b}(m) := \lambda_A(m) \quad P_{A^b} := (P_A)^* \cap \mathbb{P}_{M_A}^{\text{conc}}$$

Given a crash-aware strategy $\sigma : A \in \underline{\text{Crash}}$ we define the crash-less strategy $\sigma^b : A^b$ as below. Note that $-^b$ formalizes $\text{ops}(-)$ (§2). It is also useful to provide a reverse operation $-^\#$, read *re-crash*, that lifts, in a persistent way, a crash-less strategy $\sigma : A^b$ into a strategy $\sigma^\# : A$.

$$\sigma^b := \{\text{ops}(s) \in \mathbb{P}_{M_A}^{\text{conc}} \mid s \in \sigma\} \quad \sigma^\# := \{s \in \mathbb{P}_{M_A}^\# \mid \text{ops}(s) \in \sigma\}$$

4.4 Strict Linearizability

Similarly to how Oliveira Vale et al. [31] characterizes linearizability by lifting a non-saturated strategy to a saturated strategy, we formalize strict linearizability by lifting a strategy without crashes into a strategy with crashes.

Definition 4.6. Given games A, B , we define the strict lift $\text{str}(\sigma) : A \multimap B$ of a crash-less strategy $\sigma : A^b \multimap B^b$ as the crash-aware strategy: $\text{str}(\sigma) := K_\# \sigma^\#$

It then turns out that, similarly to what we did for crash-aware linearizability, strict linearizability supports the following refinement-based characterization:

PROPOSITION 4.7. $v' : A$ is strictly linearizable to $v : A^b$ if and only if $v' \subseteq \text{str}(v)$.

We make use of this characterization to show the following observational refinement property:

PROPOSITION 4.8 (OBSERVATIONAL REFINEMENT). Suppose $v'_A : A$ is strictly linearizable to $v_A : A^b$ and that $\sigma : A^b \multimap B^b$ implements an object linearizable to $v_B : B^b$ using v_A , i.e. $v_A; \sigma \subseteq v_B$, then, $\text{str}(\sigma)$ implements an object strictly linearizable to $\text{str}(v_B)$ using v'_A , i.e. $v'_A; \text{str}(\sigma) \subseteq \text{str}(v_B)$.

The reverse direction, unfortunately, does not hold, fundamentally because $\text{str}(\text{ccopy}_{A^b}) \neq \text{crashcopy}_{A^b}$. By similar reasoning as the locality for crash-aware linearizability, we also obtain:

PROPOSITION 4.9 (LOCALITY). For crash-aware strategies $v'_A : A, v'_B : B$ and crash-less strategies $v_A : A, v_B : B$: $v'_A \subseteq \text{str}(v_A)$ and $v'_B \subseteq \text{str}(v_B)$ if and only if $v'_A \otimes v'_B \subseteq \text{str}(v_A \otimes v_B)$

4.5 Durable Linearizability

Recall that a crash-aware play (i.e., a trace) is durable when the set of agents on different epochs is disjoint. Given a game A , let P_A^{dur} be the subset of P_A^{\ddagger} containing only its durable plays. As we noted in §2, durable plays s have the important property that their de-crash s^b is always defined. We call a crash-aware strategy *durable* if it only contains durable plays.

Now, for our refinement-based formulation, we define a durable lift $\text{dur}(-)$, which assigns to a crash-less strategy $v : A^b$ the durable strategy $\text{dur}(v) : A$ defined by $\text{dur}(v) : A := (K_{\text{Conc}} v)^{\ddagger} \cap P_A^{\text{dur}}$.

The operation $K_{\text{Conc}} -$ in the formula is defined by Oliveira Vale et al. [31] similarly to $K_{\ddagger} -$, but in the crash-less setting. It may be more intuitively understood through their result that:

$$K_{\text{Conc}} v = \{s \in P_{A^b} \mid \exists t \in v.s \rightsquigarrow t\}$$

that is to say, $K_{\text{Conc}} -$ assigns to a crash-less strategy v the smallest strategy containing v that has all plays linearizable w.r.t. to v . We observe that, indeed, $\text{dur}(-)$ does provide an appropriate lifting operation for durable linearizability.

PROPOSITION 4.10. $v' : A$ is durably linearizable to $v : A^b$ if and only if $v' \subseteq \text{dur}(v)$.

This refinement characterization enables us to show observational refinement and locality.

PROPOSITION 4.11 (OBSERVATIONAL REFINEMENT AND LOCALITY).

- Let A, B be games. Then $v'_A : A$ is durably linearizable to $v_A : A^b$ if and only if whenever $\sigma : A^b \multimap B^b$ is a crash-less strategy implementing a crash-less object linearizable to v_B using v_A , then $\text{dur}(\sigma) : A \multimap B$ implements an object durably linearizable to v_B using v'_A .
- For durable strategies $v'_A : A, v'_B : B$ and crash-less $v_A : A, v_B : B$: $v'_A \xrightarrow{\text{dur}} v_A$ and $v'_B \xrightarrow{\text{dur}} v_B$ if and only if $v'_A \otimes v'_B \xrightarrow{\text{dur}} v_A \otimes v_B$

5 PROGRAM LOGIC

In this section, we present a program logic for verifying durable linearizability, which is based on rely-guarantee reasoning, crash Hoare logic and possibility reasoning. We first (§5.1) briefly discuss how to abstract away recovery. Then (§5.2) we define an object-agnostic imperative programming language. Lastly (§5.3) we demonstrate the key rules of the program logic. We refer readers to our appendix for its variation for verifying crash-aware linearizability, which is largely similar.

5.1 Recovery

We start by discussing a simple way of removing recovery events from a play, which is enough for our purposes. First, we fix a certain kind of signature for objects with recovery.

Definition 5.1. We define a recovery signature $E \cup R$ to be the union of two effect signatures E for regular operations and R for recovery operations.

To simplify reasoning about programs with recovery, it is common to provide a way to remove the recovery events from the specification. In our setting, this is notoriously simple.

Definition 5.2. We say a strategy $v' : \dagger(E \cup R)$ recovery-refines to $v : \dagger E$ when $v' \upharpoonright_E \subseteq v$.

It is straightforward to see that the following refinement theorem holds.

PROPOSITION 5.3 (RECOVERY REFINEMENT THEOREM). *Suppose $v' : \dagger(E \cup R)$ recovery-refines to $v : \dagger E$ and that $\sigma' : \dagger(E \cup R) \multimap \dagger F$ then, for $\sigma : \dagger E \multimap \dagger F$, $\sigma' \upharpoonright_{\dagger E \multimap \dagger F} \subseteq \sigma \implies v'; \sigma' \subseteq v; \sigma$*

5.2 Programming Language

5.2.1 Syntax. We start by defining a language Com for commands over some effect signature E .

$\text{Prim} := x \leftarrow e(a) \mid \text{assume}(\phi) \mid \text{ret } v \quad \text{Com} := \text{Prim} \mid \text{Com}; \text{Com} \mid \text{Com} + \text{Com} \mid \text{Com}^* \mid \text{skip}$

Prim stands for primitive commands. The assignment command, $x \leftarrow e(a)$, executes the effect $e \in E$ with argument a and stores the response to variable x in a local environment $\Delta \in \text{Env}$. The assume command, $\text{assume}(\phi)$, takes a boolean function ϕ over Δ and terminates the computation if it evaluates to `False`. We implement loops and if-statements using $\text{assume}(-)$ in the usual way. The return command, $\text{ret } v$, stores the value v into a reserved variable `res`, and is executed once per invocation of a procedure. Com is the grammar of commands defined as usual in a Kleene algebra.

The implementation M of an object (with the effect signature $F \cup R_F$) is defined as a collection of commands $M[\alpha]^f \in \text{Com}$, $M = (M[\alpha])_{\alpha \in \Upsilon} = (M[\alpha]^f)_{\alpha \in \Upsilon, f \in F \cup R_F}$, which implements each method $f \in F \cup R_F$ per agent $\alpha \in \Upsilon$. Here F defines the overlay's regular procedures and R_F its recovery procedures. For simplicity, we require that there is only one recovery program r in R_F , i.e. $R_F = \{r : \mathbf{1} \rightarrow \mathbf{1}\}$. We call $M[\alpha]$ a local implementation and $M \in \text{CMod}$ a concurrent module, where CMod is the set of all concurrent modules.

5.2.2 Memory Model & Object State. Observe that our programming language is object-agnostic in that it operates over an arbitrary object of type E . This means that the language does not have a memory model baked in. Instead, the underlay object's effect signature E , over which the language is parameterized, determines which memory operations the user can perform. For example, to implement the FLiT memory cell in Fig. 1, one would use as the underlay a buffered memory cell with the BCell signature. Then, one can write a program with statements like $x \leftarrow B.\text{load}(); B.\text{flush}()$ to manipulate the memory shared across threads.

We define the underlay state as $(\Delta, s) \in \text{UndState}$, a tuple of a local environment Δ and a history $s \in P_{\dagger E}$. The local environment Δ is defined solely as a mapping from local variables to their values (with Δ_0 representing the empty local environment). The history s is a canonical representation for shared state, since it records all previous operations on the shared underlay object. One may reconstruct other more intuitive definitions of the shared state by defining an interpretation function over the trace s . For example, given the traces $p \in \nu_{\text{FLiT}}$ of one FLiT memory cell, we can define the evaluation function $\text{fstate} : \nu_{\text{FLiT}} \rightarrow \text{Val}$ to compute the current value of the cell by reading the latest stored value. In particular, note that we may use the (atomic) linearized specification for FLiT because of observational refinement.

5.2.3 Semantics. Primitive commands B are interpreted as state transformers $\llbracket B \rrbracket_\alpha : \text{UndState} \rightarrow \mathcal{P}(\text{UndState})$ from a set of underlay states to a new set of states. The $\llbracket B \rrbracket_\alpha$ depends on α only in that it tags event it adds to the history with an agent identifier α . We then lift the state transformer

$$\begin{array}{c}
\longrightarrow \subseteq (\text{Com} \times \text{UndState}) \times \Upsilon \times (\text{Com} \times \text{UndState}) \quad \longrightarrow_{RE} (\text{Cont} \times \text{ModState}) \times \text{CMod} \times (\text{Cont} \times \text{ModState}) \\
\frac{f \in F \quad a \in \text{par}(f) \quad \Delta' = \Delta[\alpha \mapsto [\text{arg} \mapsto a]]}{\langle c[\alpha \mapsto \text{idle}], \Delta, s \rangle \longrightarrow_{RE}^M \langle c[\alpha \mapsto M[\alpha]^f], \Delta', s \cdot \alpha : f \rangle} \text{INV} \quad \frac{\langle C, \Delta, s \upharpoonright_E \rangle \longrightarrow_{\alpha} \langle C', \Delta', s' \upharpoonright_E \rangle}{\langle c[\alpha \mapsto C], \Delta, s \rangle \longrightarrow_{RE}^M \langle c[\alpha \mapsto C'], \Delta', s' \rangle} \text{STEP} \\
\frac{\pi_{\alpha}(s \upharpoonright_F) = p \cdot f \quad \Delta(\alpha)(\text{res}) = v \in \text{ar}(f) \quad \Delta' = \Delta[\alpha \mapsto \emptyset]}{\langle c[\alpha \mapsto \text{skip}], \Delta, s \rangle \longrightarrow_{RE}^M \langle c[\alpha \mapsto \text{idle}], \Delta', s \cdot \alpha : v \rangle} \text{RET} \\
\frac{\forall \alpha \in s.c'[\alpha] = \text{dead} \quad \forall \alpha \in \Upsilon. \alpha \notin s \Rightarrow c'[\alpha] = \text{halt}}{\langle c, \Delta, s \rangle \longrightarrow_{RE}^M \langle c', \Delta_0, s \cdot \frac{1}{2} \rangle} \text{CRASH} \quad \frac{s = s' \cdot \frac{1}{2} \quad \vec{r} = \text{perm}(R_E) \quad C = \text{sequence}(\vec{r}, M[\alpha]^r)}{\langle c[\alpha \mapsto \text{halt}], \Delta, s \rangle \longrightarrow_{RE}^M \langle c[\alpha \mapsto C], \Delta, s \cdot \alpha : r \rangle} \text{STARTREC} \\
\frac{\pi_{\alpha}(s \upharpoonright_{F \cup R_F}) = s' \cdot r \quad \Delta(\alpha)(\text{res}) = v \in \text{ar}(r) \quad \Delta' = \Delta[\alpha \mapsto \emptyset] \quad \forall \alpha \in \Upsilon. c[\alpha] = \text{dead} \Rightarrow c'[\alpha] = \text{dead} \quad \forall \alpha \in \Upsilon. c[\alpha] \neq \text{dead} \Rightarrow c'[\alpha] = \text{idle}}{\langle c[\alpha \mapsto \text{skip}], \Delta, s \rangle \longrightarrow_{RE}^M \langle c', \Delta', s \cdot \alpha : v \rangle} \text{ENDREC} \\
\text{where } \text{sequence}(\vec{r}, C) = \begin{cases} C & \vec{r} = \epsilon \\ (x_r \leftarrow r(a)); \text{sequence}(\vec{r}', C) & \vec{r} = r \cdot \vec{r}' \wedge a \in \text{par}(r) \wedge \text{reserved}(x_r) \end{cases}
\end{array}$$

Fig. 5. Local Small-Step Semantics (\longrightarrow) and Module Small-step semantics (\longrightarrow_{RE})

$\llbracket B \rrbracket_{\alpha}$ to a thread-local small-step semantics $\langle C, \Delta, s \rangle \longrightarrow_{\alpha} \langle C', \Delta', s' \rangle$, which encodes how α steps through commands in a mostly standard way following the Kleene algebra structure of commands.

In Fig. 5, we lift this local small-step semantics to a concurrent module small-step semantics $\langle c, \Delta, s \rangle \longrightarrow_{RE}^M \langle c', \Delta', s' \rangle$, which takes a continuation $c \in \text{Cont} := \Upsilon \rightarrow \{\text{idle}, \text{skip}, \text{dead}, \text{halt}\} + \text{Com}$ and a module state $(\Delta, s) \in \text{ModState} := (\Upsilon \rightarrow \text{Env}) \times P_{\dagger(E \cup R_E) \rightarrow \dagger(F \cup R_F)}$ containing local environments for all agents and the global trace of the system. The first three rules come from the semantics in Oliveira Vale et al. [31] to handle mainly the execution of regular procedures:

INV Allows a new invocation of any overlay operation f in an idle thread and appends the new invocation to the end of s .

STEP Non-deterministically chooses some thread that is running a program C and performs a thread local small-step in that thread with its effect applied to the concurrent module state.

RET Allows any thread that has finished its program to return to idle by appending the return value as a response to the end of s and clearing $\Delta[\alpha]$.

We add three highlighted rules to handle crashes and recoveries:

CRASH Allows for crashes to happen at any time, resetting local environments to Δ_0 for all agents, marking all the previously active agents as dead and all remaining ones as halt.

STARTREC Non-deterministically selects a halted thread α and starts the recovery phase by using C as its continuation, which sequentially runs first a permutation of underlay recoveries ($\vec{r} = \text{perm}(R_E)$) and then the overlay recovery $M[\alpha]^r$. This is achieved by using sequence to sequence a list of commands (note that $\text{reserved}(x_r)$ simply means that x_r is a reserved variable). During the recovery phase, other threads must wait for α to finish the recovery before their executions. The execution of α follows the **STEP** rule.

ENDREC When the recovery finishes, any agent that is not dead becomes idle, so that the system can now run normally. To enforce the durable assumption, dead agents will no longer run.

We define the denotation of a module by the formula below as the set of traces generated by the small-step semantics from the initial configuration, where c_0 is the initial continuation (every agent is idle) and Δ_0 is the initial environment where every agent has an empty local environment.

$$\llbracket M \rrbracket_{RE} := \{s \mid \exists c \in \text{Cont}, \Delta \in (\Upsilon \rightarrow \text{Env}). \langle c_0, \Delta_0, \epsilon \rangle \longrightarrow_{RE}^M \langle c, \Delta, s \rangle \subseteq P_{\dagger(E \cup R_E) \rightarrow \dagger(F \cup R_F)}\}$$

5.3 A Program Logic for Durable Objects

5.3.1 Interfaces. The interface of a crash-aware linearizable object E is a (round bracket) tuple.

$$(v'_E : \dagger(E \cup R_E), v_E : \dagger E) \quad \text{s.t.} \quad v'_E \upharpoonright_E \xrightarrow{\xi} v_E$$

v'_E is the concrete specification containing all possible traces the object can produce, including crash and recovery events, and v_E is the linearized specification after removing recovery events.

Similarly, we define the interface of a durable linearizable object E as the (angle bracket) tuple but with a major difference: the durable interface's linearized specification v_E is crash-less.

$$\langle v'_E : \dagger(E \cup R_E), v_E : \dagger E \rangle \quad \text{s.t.} \quad v'_E \upharpoonright_E \stackrel{\text{dur}}{\approx} v_E$$

The objective of our program logic is to establish the judgment

$$\vdash M : (v'_E, v_E) \rightarrow (v'_F, v_F) \quad \text{or} \quad \vdash M : (v'_E, v_E) \rightarrow \langle v'_F, v_F \rangle$$

which means under the assumption that the implementation M implements F with either a crash-aware interface (v'_F, v_F) (the variation described in our appendix) or a durable interface $\langle v'_F, v_F \rangle$ (the variation we describe here), using the crash-aware underlay E with interface (v'_E, v_E) . The concrete specification v'_F is defined by running an implementation M above v'_E , i.e., $v'_F = v'_E; \llbracket M \rrbracket_{R_E}$. The program logic's soundness guarantees the validity of the crash-aware/durable overlay interface. In this context, (v'_E, v_E) is called M 's *underlay*, while $\langle v'_F, v_F \rangle$ is called M 's *overlay*.

The specifications v'_E, v_E, v'_F, v_F are fixed in the program logic. For simplicity, we take them as a parameter in all that follows and omit the parametrization in the concrete proof rules.

5.3.2 The Rely-Guarantee Crash Linearizability Hoare Logic (CLHL).

Configurations & Assertions. CLHL uses as proof configurations triples $(\Delta, s, \rho) \in \text{Config} := \text{ModState} \times \text{Poss}$, where $\rho \in \text{Poss}$, called a possibility, is a play of type $\dagger F$ linearizable w.r.t. v_F . A configuration is valid when s is durably linearizable to ρ and ρ is linearizable w.r.t. v_F . This ensures that the concrete trace s is always durably linearizable with respect to v_F after the recovery refinement. Pre-conditions P , post-conditions Q , and crash conditions Q_i are given by sets of configurations, while rely conditions \mathcal{R} and guarantee conditions \mathcal{G} are relations over Config.

Top Level Rules. The top level rule OBJECT IMPL proves that M implements the overlay $\langle v'_F, v_F \rangle$ using the underlay (v'_E, v_E) . It requires the prover to find an object invariant $I : Y \rightarrow \mathcal{P}(\text{Config})$ for the implementation and then verify regular procedures and the recovery separately.

$$\frac{\begin{array}{l} \forall \alpha, \alpha' \in Y. \alpha \neq \alpha' \Rightarrow \mathcal{G}[\alpha] \cup \text{invoke}_\alpha(-) \cup \text{return}_\alpha(-) \subseteq \mathcal{R}[\alpha'] \\ \forall \alpha \in Y. \mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \models_\alpha^F M[\alpha] \quad \forall \alpha \in Y. I \models_\alpha^R M[\alpha] \end{array}}{\vdash M : (v'_E, v_E) \rightarrow \langle v'_F, v_F \rangle} \quad \text{OBJECT IMPL}$$

Verifying Regular Procedures. To verify a concurrent object, the OBJECT IMPL rule requires finding appropriate rely \mathcal{R} and guarantee \mathcal{G} for the object. The rely $\mathcal{R}[\alpha']$ of an agent models the interference of other threads in the executions and therefore must take into account at least invocations, returns, and the guarantee of other agents α (specified, respectively, by $\text{invoke}_\alpha(-)$, $\text{return}_\alpha(-)$, and $\mathcal{G}[\alpha]$). The prover needs to show $\mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \models_\alpha^F M[\alpha]$, which asserts that when α runs regular methods in F , assuming other threads behave according to $\mathcal{R}[\alpha]$, α will behave according to $\mathcal{G}[\alpha]$, and $I[\alpha]$ is satisfied when the thread α is idle.

The LOCAL IMPL rule proves this judgment by splitting $I[\alpha]$ into conjunctions of $P[\alpha]^f$, each specifying the pre-condition of a method invocation, and then proving a series of objectives ($- \circ -$ stands for relational composition).

$$\frac{\begin{array}{l} I[\alpha] = \bigcap_{f \in F} P[\alpha]^f \quad \forall f \in F. (\Delta_0, \epsilon, \epsilon) \in P[\alpha]^f \quad \forall f \in F. \text{stable}(\mathcal{R}[\alpha], P[\alpha]^f) \\ \forall f \in F. \mathcal{R}[\alpha], \mathcal{G}[\alpha] \models_\alpha^f \{P[\alpha]^f\} M[\alpha]^f \{Q[\alpha]^f\} \{\top\} \quad \forall f \in F. \text{return}_\alpha(f) \circ Q[\alpha]^f \subseteq I[\alpha] \end{array}}{\mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \models_\alpha M_F[\alpha]} \quad \text{LOCAL IMPL}$$

Firstly, each pre-condition $P[\alpha]^f$ needs to include the initial configuration and must be stable under interferences $\mathcal{R}[\alpha]$ of the environment, which implies the invariant $I[\alpha]$ to be stable.

Then, the prover needs to show that $\mathcal{R}[\alpha], \mathcal{G}[\alpha] \vDash_{\alpha}^f \{P[\alpha]^f\} M[\alpha]^f \{Q[\alpha]^f\} \{\top\}$ is satisfied for each method f . The hexad $\mathcal{R}, \mathcal{G} \vDash_{\alpha}^f \{P\} C \{Q\} \{Q_{\downarrow}\}$ means that given states satisfying P , running the program C on thread α in an environment with interference in \mathcal{R} will produce actions in \mathcal{G} , and if it terminates normally, the state will satisfy Q , and if it crashes, the state will satisfy Q_{\downarrow} . A hexad is proved with proof rules introduced later. It is worth mentioning that *there is no need to explicitly specify and prove a crash condition for any regular method*, and we can simply put \top as the crash condition. This is true because:

- (1) The guarantee $\mathcal{G}[\alpha]$ of the current thread is included in any other thread's rely $\mathcal{R}[\alpha']$, and therefore any step during the execution of any method in thread α is captured in $\mathcal{R}[\alpha']$.
- (2) For any other thread α' , its invariant $I[\alpha']$ is stable w.r.t. $\mathcal{R}[\alpha']$, which means any state after any execution step of any method in thread α (captured in $\mathcal{R}[\alpha']$) is in $I[\alpha']$.
- (3) Therefore, the state of thread α will satisfy any other thread's invariant $I[\alpha']$ at any time (including the point of crash), and the crash condition in α can be derived from $I[\alpha']$.

Lastly, after finishing the execution of a method and returning from it, the invariant $I[\alpha]$ needs to be satisfied so that the current thread can still access the object by invoking its procedures.

Verifying Recovery. Then, to ensure the durability of the object, provers need to show $I \vDash_{\alpha}^R M[\alpha]$, which means whenever a crash happens, the execution of the recovery on any thread α can restore the program state to satisfy the object invariant I . It can be verified via the RECOVER IMPL rule.

$$\frac{\text{ID}, \top \vDash_{\alpha}^r \{P_r[\alpha]\} M[\alpha]^r \{Q_r[\alpha]\} \{Q_{\downarrow}[\alpha]\} \quad Q_{\downarrow}[\alpha] \subseteq P_r[\alpha] \quad \cup_{\alpha' \in \Upsilon} I[\alpha'] \Rightarrow_{\downarrow} Q_{\downarrow}[\alpha] \quad \text{return}_{\alpha}(r) \circ Q_r[\alpha] \subseteq \cap_{\alpha' \in \Upsilon} I[\alpha']}{I \vDash_{\alpha}^R M[\alpha]} \text{RECOVER IMPL}$$

The prover needs to find a recovery pre-condition P_r , a recovery post-condition Q_r , and a crash condition Q_{\downarrow} for the recovery program, and prove $\text{ID}, \top \vDash_{\alpha}^r \{P_r[\alpha]\} M[\alpha]^r \{Q_r[\alpha]\} \{Q_{\downarrow}[\alpha]\}$, which means running the recovery program $M[\alpha]^r$ from states in $P_r[\alpha]$ will either recover the system into states in $Q_r[\alpha]$ or crash into states in $Q_{\downarrow}[\alpha]$. Since the recovery program always runs after a crash, the crash condition Q_{\downarrow} needs to imply P_r . But as the recovery program executes sequentially, with no interference from other threads, the rely and guarantee for it are ID and \top .

The invariant $I[\alpha']$ serves as the crash condition of other threads. Therefore, we require that all $I[\alpha']$ crash into the crash condition Q_{\downarrow} of the recovery program. The crash-into relation (\Rightarrow_{\downarrow}) amounts to implication after adding a crash: $I \Rightarrow_{\downarrow} Q_{\downarrow} \iff \forall (\Delta, s, \rho) \in I. (\Delta_0, s \cdot \downarrow, \rho) \in Q_{\downarrow}$.

Lastly, after the execution of the recovery, the system is restored and ready to run, so the program state after the recovery's return needs to imply the invariant $I[\alpha']$ of any thread α' .

The Core Proof Rule. According to these top-level rules, proofs of both the regular procedures and the recovery boil down to proofs of hexads like $\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\} C \{Q\} \{Q_{\downarrow}\}$. Among CLHL proof rules for the hexad, the core proof rule for proving the durable linearizability is the PRIM rule, which we focus on in this section and refer readers to the appendix for other proof rules, which are standard.

$$\frac{P \Rightarrow_{\downarrow} Q_{\downarrow} \quad Q \Rightarrow_{\downarrow} Q_{\downarrow} \quad Q_{\downarrow} \Rightarrow_{\downarrow} Q_{\downarrow} \quad \text{stable}(\mathcal{R}, P) \quad \text{stable}(\mathcal{R}, Q) \quad \mathcal{G} \vdash_{\alpha} \{P\} B \{Q\}}{\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\} B \{Q\} \{Q_{\downarrow}\}} \text{PRIM}$$

There are three groups of PRIM rule's premises. Firstly, as crashes can happen at any point, the pre-/post-condition and the crash condition should be able to crash into (\Rightarrow_{\downarrow}) the crash condition. Then, as any rely-guarantee logic, the pre-/post-condition needs to be stable w.r.t. the rely \mathcal{R} .

$$\begin{aligned} \mathcal{G} \vdash_{\alpha} \{P\} B \{Q\} &\iff \forall \Delta, s, \rho, \Delta', s'. ((\Delta, s, \rho) \in P \wedge (\Delta', s') \in \llbracket B \rrbracket_{\alpha}(\Delta, s) \cap \nu E) \quad \text{where } \rho \dashrightarrow \rho' \iff \\ &\implies (\exists \rho'. (\Delta', s', \rho') \in Q \wedge (\Delta, s, \rho) \mathcal{G}(\Delta', s', \rho') \wedge \rho \dashrightarrow \rho') \quad \exists t_P \in (M_P^P)^*. \rho \cdot t_P \rightsquigarrow_{\vdash F} \rho' \end{aligned}$$

Lastly, we need to prove the commit rule $\mathcal{G} \vdash_{\alpha} \{P\} B \{Q\}$ for the primitive command B . It states that after a step from a state in P made by the command B the new state will satisfy the post-condition Q

and the guarantee. This step may be the commitment point of some pending operations. To maintain the invariant that s is durably linearizable to ρ , the commit rule allows an angelic linearization update, $\rho \dashrightarrow \rho'$, where provers can append several response events to ρ and rewrite it according to $\rightsquigarrow_{\dagger F}$ to obtain ρ' , a new possibility that s linearizes into. Moreover, since possibility updates are recorded in \mathcal{G} , its effect is visible to any other thread. Only a careful choice of possibility updates will respect other threads' relies and prove that this object is indeed durably linearizable.

Soundness. CLHL is justified by the following soundness theorem.

PROPOSITION 5.4 (SOUNDNESS). *If $\vdash M : (v'_E, v_E) \rightarrow \langle v'_F, v_F \rangle$ is provable, and (v'_E, v_E) is a valid crash-aware interface, and $v'_F = v'_E; \llbracket M \rrbracket_{R_E}$, then $\langle v'_F, v_F \rangle$ is a valid durable interface.*

5.4 Examples Revisited

In this section, we present some high-level proof ideas of the FLiT example and demonstrate the usage of the program logic. The FLiT object is built above the buffered memory cell BCell.

The Buffered Cell. The buffered memory cell's concrete traces in v'_{BCell} are crash-aware linearizable to its specification v_{BCell} , which we can define through an interpretation function, $\text{mstate} : v_{\text{BCell}} \rightarrow \mathcal{P}(\text{Val} \times \text{Val})$, which computes the set of all possible combinations of the persisted value (the first component) and the buffered value (the second component), as seen in §2.3.

Using mstate , the specification v_{BCell} is essentially defined as the set of traces that can step from the initial state, the singleton set $\{(v_0, v_0)\}$, to some non-empty state, with the step function below. The sets on the two sides of the arrow are the value of mstate before and after appending the events to the trace.

$$S \xrightarrow{\alpha:\text{store}(v) \cdot \alpha:\text{ok}} \{(v_p, v) \mid (v_p, v_b) \in S\} \cup \{(v, v)\} \qquad S \xrightarrow{\alpha:\text{flush} \cdot \alpha:\text{ok}} \{(v_b, v_b) \mid (v_p, v_b) \in S\}$$

$$S \xrightarrow{\zeta} \{(v_p, v_p) \mid (v_p, v_b) \in S\} \qquad S \xrightarrow{\alpha:\text{load} \cdot \alpha:\text{ok}(v)} \{(v_p, v) \mid (v_p, v) \in S\}$$

- When a store operation finishes, there are two possible outcomes: the value may have been stored only to the buffered content, while the persisted content remains the same as before the store; the value may be persisted, making the buffered content the same as the persisted one.
- When a flush operation finishes, the buffered value gets flushed into the persisted part. Since after each store operation, the buffered content is uniquely determined (synchronized), after a consequent flush operation, the content of mstate is uniquely determined.
- When a crash ζ happens, the buffered content is lost, and after the crash, the buffered content is overwritten by the persisted value, which may have various possibilities because a flush may not have happened before the crash. As a result, the uniqueness of the buffered content no longer holds after the crash and is un-synchronized. *The non-determinism brought by store and ζ is the first challenge of the FLiT proof and the reason we define mstate in this way.*
- When a load operation finishes, the actual buffered content is determined and all future load will not observe other possibilities of the buffered content. As we will explain later, this behavior makes the load operation an external linearization point of buffered operations before a crash. *The helping mechanism, especially helps across crashes, is the second challenge of the FLiT proof.* The returned value must be consistent with at least one possible buffered content in S . Otherwise, the post-state is an empty set and this trace will not be accepted in v_{BCell} .

To use CLHL to verify the FLiT overlay, we need an invariant I that links the overlay and underlay states and is maintained by any program step. Depending on the current buffered memory cell state, we split the invariant into three cases. (1) When the buffered content v_b is synchronized and persisted (the Flushed state), then the overlay state $\text{fstate}(\rho)$ should also be v_b , i.e., the store operation that writes this v_b is durably linearized. (2) When the buffered content v_b is synchronized

but not persisted (the Unflushed state), we use a ghost list B to buffer the pending overlay store(v_b) operations in order, so future operations can help linearize it when the value gets persisted. (3) When a crash happens (the Unsynced state), the buffered content v_b is un-synchronized and corresponds to some store(v_b) operation in the ghost list B , in case it has persisted, or is equal to the current overlay state $\text{fstate}(\rho)$, when none of the buffered operations persisted.

As a result, the proof configuration now becomes $(\Delta, s, \rho, B) \in \text{ModState} \times \text{Poss} \times M_F^*$. According to the OBJECT IMPL rule, we need to find rely and guarantee conditions verifying $\mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \models_{\alpha}^F M[\alpha]$ for the load and store operations and $I \models_{\alpha}^R M[\alpha]$ for an empty recovery procedure.

5.4.1 Regular Procedure Proofs. To prove regular procedures through the LOCAL IMPL rule, we must find the pre-/post-conditions corresponding to each procedure and prove their Hoare quadruples. For the FLIT implementation, we prove Hoare quadruple (1) and (2) for the load and store operations.

$$\mathcal{R}[\alpha], \mathcal{G}[\alpha] \models_{\alpha} \{\text{invoke}_{\alpha}(\text{load}) \circ I\} \text{load}() \{\text{returned}_{\alpha}(\text{load}) \circ I\} \{\top\} \quad (1)$$

$$\mathcal{R}[\alpha], \mathcal{G}[\alpha] \models_{\alpha} \{\text{invoke}_{\alpha}(\text{store}) \circ I\} \text{store}() \{\text{returned}_{\alpha}(\text{store}) \circ I\} \{\top\} \quad (2)$$

The invoke and returned relations are defined below. The invoke simply adds an invocation (by clients of the overlay object) to the procedure f to the end of s and ρ . The returned asserts the returned result recorded in Δ is consistent with the one linearized in ρ by the prover.

$$\begin{aligned} (\Delta, s, \rho) \text{invoke}_{\alpha}(f)(\Delta', s', \rho') &\iff \left(\begin{array}{l} (\Delta, s, \rho) \in \text{idle}_{\alpha} \wedge \exists a. \Delta'(\alpha) = [\text{arg} \mapsto a] \wedge \\ \forall \alpha' \neq \alpha. \Delta'(\alpha') = \Delta(\alpha) \wedge s' = s \cdot \alpha : f \wedge \rho' = \rho \cdot \alpha : f \end{array} \right) \\ (\Delta, s, \rho) \text{returned}_{\alpha}(f)(\Delta', s', \rho') &\iff (\Delta', s', \rho') = (\Delta, s, \rho) \wedge \exists v \in \text{ar}(f). \Delta(\alpha)(\text{ret}) = v \wedge \text{last}(\pi_{\alpha}(\rho)) = \alpha : v \end{aligned}$$

These quadruples are proved by mainly using the PRIM rule to step through primitive commands. In most of the cases, the underlay load/store operations only add pending overlay operations to the list B , and a consequent flush operation makes sure they are persisted and helps operations in B linearize. The Counter object prevents unnecessary flushes in this process but is not the main complexity of the FLIT object, and thus we refer readers to the appendix for its treatment.

Figure 6 shows the proof outline for the load operation, which we use as an example for demonstration. The program contains two potential linearization points, line 2 and line 5, and we show how to use the PRIM rule to complete proofs and find linearizations at these points.

The underlay load operation at line 2 may execute from three different situations depending on the object state (Flushed, Unflushed, Unsynced). We choose to perform three different updates to the possibility ρ and the ghost list B and illustrate them through guarantee conditions below, which record the effects of these updates on proof configurations.

$$(s, \rho, B) \mathcal{G}_{\text{load-f}}[\alpha](s', \rho', B') \iff \left(\begin{array}{l} \exists v. \text{Flushed}(s, B) \wedge (v, v) \in \text{mstate}(s \uparrow_{\text{BCell}^i}) \wedge v = \text{fstate}(\rho) \wedge \\ \text{lin}(\rho') = \text{lin}(\rho) \cdot \alpha : \text{load} \cdot \alpha : v \wedge B' = \epsilon \wedge s' = s \cdot \alpha : M.\text{load} \cdot \alpha : v \end{array} \right) \quad (3)$$

$$(s, \rho, B) \mathcal{G}_{\text{load-uf}}[\alpha](s', \rho', B') \iff \left(\begin{array}{l} \exists v. \text{Unflushed}(s, B) \wedge \text{last}(B \uparrow_{\text{store}}) = \text{store}(v) \wedge \\ B' = B \cdot \alpha : \text{load} \wedge \rho' = \rho \wedge s' = s \cdot \alpha : M.\text{load} \cdot \alpha : v \end{array} \right) \quad (4)$$

$$(s, \rho, B) \mathcal{G}_{\text{load-us}}[\alpha](s', \rho', B') \iff \left(\begin{array}{l} \exists v, B_1, B_2. \text{Unsynced}(s, B) \wedge (v, v) \in \text{mstate}(s \uparrow_{\text{BCell}^i}) \wedge \\ s' = s \cdot \alpha : M.\text{load} \cdot \alpha : v \wedge \left(\begin{array}{l} B = B_1 \cdot B_2 \wedge \text{last}(B_1) = \text{store}(v) \\ \vee (\text{fstate}(\rho) = v \wedge B_1 = \epsilon) \end{array} \right) \wedge \\ \text{lin}(\rho') = \text{merge}(\text{lin}(\rho), B_1) \cdot \alpha : \text{load} \cdot \alpha : v \wedge B' = \epsilon \end{array} \right) \quad (5)$$

Load from Flushed State. When the underlay memory cell is at the Flushed state, i.e., there are no buffered operations and $B = \epsilon$, then, the current memory content $\text{fstate}(\rho)$ is exactly the same

```

{invokeα(load) ◦ I}
1: load() {
  {I ∧ α:load ∈ sO ∧ (Flushed ∨ Unflushed ∨ Unsynced)} // Pload
2: v ← M.load(); // load-f/load-uf/load-us
  {I ∧ ((Flushed ∧ last(πα(ρ)) = v) ∨
  (Unflushed ∧ (∃B'.B' · α':store(v) · α:load ⊆ B ∨ last(πα(ρ)) = v)))} // Qload
3: n ← C.get();
  {I ∧ ((n = 0 ∧ last(πα(ρ)) = v) ∨
  (n ≠ 0 ∧ (∃B'.B' · α':store(v) · α:load ⊆ B ∨ last(πα(ρ)) = v))) ∧ (Flushed ∨ Unflushed)}
4: if(n ≠ 0) {
  {I ∧ (Flushed ∨ Unflushed) ∧ (∃B'.B' · α':store(v) · α:load ⊆ B ∨ last(πα(ρ)) = v)}
5:   M.flush(); // flush
  {I ∧ last(πα(ρ)) = v}
6: }
  {I ∧ last(πα(ρ)) = v}
7: ret v
8: }
{returnedα(load) ◦ I} {T}

```

Fig. 6. A Proof Snippet of the load operation of FLiT Memory Cell

as the content in the underlay memory cell v . Therefore, we can simply extend the linearized prefix $\text{lin}(\rho)$ in ρ with $\alpha:\text{load} \cdot \alpha:v$ by reordering the pending load to the place and add the response as (3).

Load from Unflushed State. When the underlay memory is at the Unflushed state, there are different possible values for the persisted content. Although the underlay load will load the most recently buffered value v , we do not know whether v has been persisted or not. If a crash happens before returning from the current overlay load, this value may be lost from the memory and we are not supposed to linearize $\alpha:\text{load} \cdot \alpha:v$ to $\text{lin}(\rho)$. Therefore, instead of linearizing it at this point, we choose to append the pending load to the buffered list B so that a subsequent flush operation from either the current program or other threads can help linearize it as (4).

Load from Unsynced State. The most special case is when the load is executed after a crash with some buffered store not flushed yet. As explained before, both the buffered and the persisted contents may have various values depending on previously buffered stores. The load operation will determine the actual content in the memory cell, which reveals and linearizes the operations that are persisted before the crash, making it an *external linearization point across crashes*.

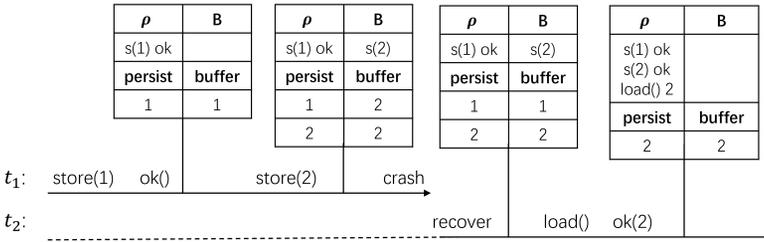


Fig. 7. External Linearization Point and Crash: the tables above the timeline show the content of the linearized trace ρ and the ghost list B in the first two rows and the mstate content in the remaining rows. $s(-)$ is a shorthand for the store $(-)$ operation.

Figure 7 shows an example of this kind of load operation. After a buffered store(2) operation, the persisted data has not been synchronized with the buffered value 2 since no flush has been performed, and the system crashes at this moment, resulting in a state with unknown content of the buffered cell. Just like (4), buffered store operations will be put into the list B instead of directly linearized into ρ . If the result of the load operation following the recovery is 2, like in this example, it implies that the buffered store operation has been persisted before the crash, and thus we can linearize the store(2) cached in B followed by the current load operation. In the other case, where the load after recovery gets 1, we know the buffered store operation failed to persist, and thus we do not linearize the store(2) and instead remove it from the list B .

We follow this pattern and modify the proof configuration as (5). We maintain as an object invariant that any persisted value in the underlay memory corresponds to some store in B or $\text{lin}(\rho)$. Based on the return value v of the underlay load, we decide how to handle buffered operations in B . If v is the result of some store(v) in B , then we know this store(v) has persisted before the crash, and we linearize all operations B_1 (by reordering them before the crash, adding responses to invocations in B_1 and putting them after their corresponding invocations) in B preceding this store into $\text{lin}(\rho)$ along with the current load operation and discard what remains in B .

Then by merging these three branches into one Hoare quadruple through the disjunction rule and weakening the post-condition to the stable Q_{load} , we prove the Hoare quadruple

$$\mathcal{R}[\alpha], \mathcal{G}[\alpha] \models_{\alpha} \{P_{\text{load}}\}v \leftarrow M.\text{load}()\{Q_{\text{load}}\}\{\top\}$$

at line 2 in Figure 6. According to the PRIM rule, the quadruple is provable because we can prove $\mathcal{G} \vdash_{\alpha} \{P_{\text{load}}\}v \leftarrow M.\text{load}()\{Q_{\text{load}}\}$ by our reasoning in previous paragraphs, i.e., any update obeys the rewrite relation $\rightsquigarrow_{\text{F}}$, and other entailments and stability checks are all true.

The post-condition Q_{load} indicates that either the current load is linearized and it is obvious that the returned value v is equal to the linearized value v , or the state is unflushed and the current load is buffered in B . In the second case, the proof that remains to be done for the rest of commands is still non-trivial. Specifically, the current load may be linearized by some external operations, or it will be linearized when the flush at line 5 takes place and we need to prove it is a valid linearization step. The proof of either case will follow the outline in Figure 6 and we can prove (1). We can also prove (2) and we refer readers to the appendix for its detailed proof.

5.4.2 Recovery Procedure Proof. The FLiT object has no recovery procedure, and therefore we use the empty recovery signature $R_{\emptyset} := \{r_{\emptyset}\}$ with the recovery program, $M[\alpha]^{r_{\emptyset}} := r() \{ \text{ret ok} \}$. According to the RECOVER IMPL rule, we need to prove the hexad ID, $\top \models_{\alpha} \{I\}M[\alpha]^{r_{\emptyset}}\{I\}\{I\}$ for r_{\emptyset} , which reduces to the idempotence of the invariant w.r.t. crashes, i.e., $I \Rightarrow_{\text{c}} I$.

As we have shown $\mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \models_{\alpha}^F M[\alpha]$ and $I \models_{\alpha}^R M[\alpha]$ for any $\alpha \in \Upsilon$, according to the OBJECT IMPL rule, we prove $\vdash M_{\text{FLiT}} : (v'_{\text{BCell}} \otimes v'_{\text{Counter}}, v_{\text{BCell}} \otimes v_{\text{Counter}}) \rightarrow (v'_{\text{FLiT}}, v_{\text{FLiT}})$, i.e., the FLiT memory cell is durably linearizable. Based on the FLiT memory cell, we implement a durable version of the one-shot write-snapshot object [6], a famous interval-sequential [7] concurrent object. We prove its linearizability using the logic in Oliveira Vale et al. [31] and use the FLiT correctness theorem 1.1 to derive its durable linearizability.

We also prove the transactional file system to be crash-aware linearizable with the crash-aware linearizability variant of CLHL. It demonstrates CLHL's ability to verify non-trivial recoveries, and to decompose complicated systems into multiple layers with simpler proofs and then easily compose these proofs to obtain the originally challenging proof of the entire system.

6 RELATED WORKS

Game Semantics. Our game semantics model is directly based on that of Ghica and Murawski [16], Oliveira Vale et al. [31], and our use of object-based game semantics traces back to Oliveira Vale et al. [30], Reddy [35, 36]. To develop our crash-aware model, we indirectly made use of insights from Mellies [28]. In its goal of describing systems written in imperative languages, our game semantics is related to some of the work by Ghica and Tzevelekos [17], Koenig and Shao [25]. It is important to note that crashes are not accurately modeled as a separate computational agent responsible for issuing crashes: crashes are instantaneous and pervasive, synchronous across components, are not invoked, and are unimplemented. Because of this, our crash-aware model is rather unorthodox in that it breaks the tradition of having only two players (Opponent and Proponent) by adding an extra player for crash events. It models crash events differently from usual moves in traditional game models by having crash events happen instantaneously and synchronously across all components, while typically, a move belongs to a single component and happens mostly asynchronously. As far as we are aware, this is the first game semantics of its kind. Because of this, while we build on the model from Oliveira Vale et al. [31] and benefit significantly from the theory there, our model needs to address the intrusive effects of properly modeling crashes.

Linearizability with Crashes. We already discussed some of the history of linearizability criteria with crashes throughout the paper [2, 4, 19, 22]. In our paper, we address strict linearizability (in the context of full-system crashes) and durable linearizability. We generalize both of them by not requiring the linearized specifications to be atomic and by allowing for blocking objects. This makes our variations of these linearizability criteria closer to interval-sequential linearizability [7]. We formulate these criteria in the style of compositional linearizability [31], which is novel. This allows us to give simple proofs of locality, develop a compositional verification framework around these criteria, give the first proof of observational refinement properties for these two criteria, and provide a counterpart to the analogous result proved for Herlihy-Wing linearizability [14] and for compositional linearizability [31]. We also discover that the inherent notion of linearizability to crash-aware objects is the linearizability criterion we called crash-aware linearizability (§4) satisfying locality and observational refinement. Although related to strict linearizability, it does not appear elsewhere. We note that while crash-aware linearizability is the compositional linearizability [31] one gets from our model **Crash**, our formulations of strict and durable linearizability impose new challenges and new structures, in particular, because they relate two distinct models of computation (concurrency with and without crashes). We conjecture that this different structure can be reconciled with that from compositional linearizability through a weakening of the notion of a Grothendieck fibration, following ideas from functorial refinement [29].

Verification with Crashes. There are approaches for verifying systems with crashes that do not involve linearizability. Much of the work on this line has been done in the context of file system verification. A perhaps notable start is the development of Crash Hoare Logic [11], later refined into recovery refinement [8], and generalized to handle concurrent systems [9, 10]. Of these, only Chajed et al. [8], which only handles sequential systems, formally proves a refinement theorem that enables building large systems. The later variants that handle concurrency lack such a contextual refinement theorem. These works, different from ours, have been mechanized.

Another important work is Khyzha and Lahav [24], which proves a contextual refinement theorem for programs with crashes. Quite interesting is the fact that their approach is reminiscent of that used by Oliveira Vale et al. [31] and by us, in that they define a notion of refinement by composition with a “Most General Client”. This most general client seems to be a special case of the copycat strategies that appear in our game models. Since they do this using operational

semantics, we believe their work is further evidence of the practicality of our approach. Moreover, their programming language features a buffered memory interface with global flushes, which our example does not. Despite the similarities, they only address linearizability by providing a few examples where linearizability specifications can be encoded in their framework, but they do not describe a generic framework to do so, nor prove a formal connection with linearizability. Modeling a memory model with global flushes in our model is straightforward: its specification is almost the same as our buffered memory cell arrays, but with a requirement of proving a memory separation property, like they had to do. We do not do this here as it was not required for our examples.

A recent line of work proves linearizability specifications, but only for a single component [13], and focuses on data structures implemented on top of NVM only. It is quite impressive in that it assumes a weak memory model, which requires handling weak consistency models, which we do not. Despite that, they do not provide a program logic and are closer to axiomatic approaches, which could hinder scalability. It is unlikely that their framework could be generalized to a compositional verification methodology without significant effort.

Concurrently to our work Bodenmüller et al. [5] verified the FLiT library and have a mechanized proof of correctness. Part of their simulation-based technique is reminiscent of our use of refinement and $\text{dur}(-)$, which they define as a specific transformation of a state-transition system into another and do not note its relationship to the structure of some compositional model (which they do not develop). Their technique is restricted to durable linearizability w.r.t. atomic specifications and is specialized in verifying persistency libraries over NVM. Our work is, therefore, significantly more general in scope. Our FLiT correctness theorem shows that linearizable objects in the sense of Oliveira Vale et al. [31] are transformed into durable linearizable libraries in our sense, and therefore applies even to non-atomic and blocking objects, proving a stronger correctness theorem for FLiT (in fact, stronger than the FLiT author's informal claim of correctness, for the same reasons).

Our program logic is the first to verify a linearizability criterion with crashes. It is based on Khyzha et al. [23], Oliveira Vale et al. [31], and takes inspiration from Crash Hoare Logic and Argosy [8]. It differs from the aforementioned works in that it proves durable, and crash-aware linearizability specifications. The compositional framework, which we directly connect with our program logic, is the only one that simultaneously provides refinement, linearizability specifications, and vertical and horizontal composition. Our theory allows us to state the correctness of systems like FLiT [38]. We also show we can verify a simplified variant of a file system interface. Note that previous file system interfaces are not verified against linearizability specifications, which are deemed as more intuitive than the kind of specifications one gets from HOCAP style specifications [12, 37].

ACKNOWLEDGMENTS

We thank the reviewers of the current and previous iterations of this work for their thoughtful revisions. This material is based upon work supported in part by NSF grants 2313433, 2019285, and 1763399, and by the Defense Advanced Research Projects Agency (DARPA) and Naval Information Warfare Center Pacific (NIWC Pacific) under Contract No. N66001-21-C-4018. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

- [1] Samson Abramsky and Guy McCusker. 1999. Game Semantics. In *Computational Logic*, Ulrich Berger and Helmut Schwichtenberg (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–55. https://doi.org/10.1007/978-3-642-58622-4_1
- [2] Marcos K Aguilera and Svend Frølund. 2003. Strict linearizability and the power of aborting. *Technical Report HPL-2003-241* (2003).

- [3] Naama Ben-David, Michal Friedman, and Yuanhao Wei. 2022. Brief Announcement: Survey of Persistent Memory Correctness Conditions. In *36th International Symposium on Distributed Computing (DISC 2022) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 246)*, Christian Scheideler (Ed.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 41:1–41:4. <https://doi.org/10.4230/LIPIcs.DISC.2022.41>
- [4] Ryan Berryhill, Wojciech Golab, and Mahesh Tripunitara. 2016. Robust Shared Objects for Non-Volatile Main Memory. In *19th International Conference on Principles of Distributed Systems (OPODIS 2015) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 46)*, Emmanuelle Anceaume, Christian Cachin, and Maria Potop-Butucaru (Eds.). Schloss Dagstuhl, Dagstuhl, Germany, 20:1–20:17. <https://doi.org/10.4230/LIPIcs.OPODIS.2015.20>
- [5] Stefan Bodenmüller, John Derrick, Brijesh Dongol, Gerhard Schellhorn, and Heike Wehrheim. 2024. A Fully Verified Persistency Library. In *Verification, Model Checking, and Abstract Interpretation*, Rayna Dimitrova, Ori Lahav, and Sebastian Wolff (Eds.). Springer Nature Switzerland, Cham, 26–47. https://doi.org/10.1007/978-3-031-50521-8_2
- [6] Elizabeth Borowsky and Eli Gafni. 1993. Immediate atomic snapshots and fast renaming. In *Proceedings of the Twelfth Annual ACM Symposium on Principles of Distributed Computing (Ithaca, New York, USA) (PODC '93)*. Association for Computing Machinery, New York, NY, USA, 41–51. <https://doi.org/10.1145/164051.164056>
- [7] Armando Castañeda, Sergio Rajsbbaum, and Michel Raynal. 2015. Specifying Concurrent Problems: Beyond Linearizability and up to Tasks. In *Proceedings of the 29th International Symposium on Distributed Computing - Volume 9363 (Tokyo, Japan) (DISC 2015)*. Springer-Verlag, Berlin, Heidelberg, 420–435. https://doi.org/10.1007/978-3-662-48653-5_28
- [8] Tej Chajed, Joseph Tassarotti, M. Frans Kaashoek, and Nikolai Zeldovich. 2019. Argosy: verifying layered storage systems with recovery refinement. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (Phoenix, AZ, USA) (PLDI 2019)*. Association for Computing Machinery, New York, NY, USA, 1054–1068. <https://doi.org/10.1145/3314221.3314585>
- [9] Tej Chajed, Joseph Tassarotti, M. Frans Kaashoek, and Nikolai Zeldovich. 2019. Verifying concurrent, crash-safe systems with Perennial. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles (Huntsville, Ontario, Canada) (SOSP '19)*. Association for Computing Machinery, New York, NY, USA, 243–258. <https://doi.org/10.1145/3341301.3359632>
- [10] Tej Chajed, Joseph Tassarotti, Mark Theng, Ralf Jung, M. Frans Kaashoek, and Nikolai Zeldovich. 2021. GoJournal: a verified, concurrent, crash-safe journaling system. In *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21)*. USENIX Association, 423–439.
- [11] Haogang Chen, Daniel Ziegler, Tej Chajed, Adam Chlipala, M. Frans Kaashoek, and Nikolai Zeldovich. 2015. Using Crash Hoare logic for certifying the FSCQ file system. In *Proceedings of the 25th Symposium on Operating Systems Principles (Monterey, California) (SOSP '15)*. Association for Computing Machinery, New York, NY, USA, 18–37. <https://doi.org/10.1145/2815400.2815402>
- [12] Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J. Parkinson, and Viktor Vafeiadis. 2010. Concurrent Abstract Predicates. In *ECOOP 2010 – Object-Oriented Programming*, Theo D'Hondt (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 504–528. https://doi.org/10.1007/978-3-642-14107-2_24
- [13] Emanuele D'Osualdo, Azalea Raad, and Viktor Vafeiadis. 2023. The Path to Durable Linearizability. *Proc. ACM Program. Lang.* 7, POPL, Article 26 (jan 2023), 27 pages. <https://doi.org/10.1145/3571219>
- [14] Ivana Filipovic, Peter O'Hearn, Noam Rinetzky, and Hongseok Yang. 2010. Abstraction for Concurrent Objects. *Theor. Comput. Sci.* 411, 51–52 (dec 2010), 4379–4398. <https://doi.org/10.1016/j.tcs.2010.09.021>
- [15] Dan R. Ghica. 2019. The far side of the cube. *CoRR abs/1908.04291* (2019). arXiv:1908.04291 <http://arxiv.org/abs/1908.04291>
- [16] Dan R. Ghica and Andrzej S. Murawski. 2004. Angelic Semantics of Fine-Grained Concurrency. In *Foundations of Software Science and Computation Structures*, Igor Walukiewicz (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 211–225. <https://doi.org/10.1016/j.apal.2007.10.005>
- [17] Dan R. Ghica and Nikos Tzevelekos. 2012. A System-Level Game Semantics. *Electronic Notes in Theoretical Computer Science* 286 (2012), 191–211. <https://doi.org/10.1016/j.entcs.2012.08.013> Proceedings of the 28th Conference on the Mathematical Foundations of Programming Semantics (MFPS XXVIII).
- [18] Éric Goubault, Jérémy Ledent, and Samuel Mimram. 2018. Concurrent Specifications Beyond Linearizability. In *22nd International Conference on Principles of Distributed Systems (OPODIS 2018) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 125)*, Jiannong Cao, Faith Ellen, Luis Rodrigues, and Bernardo Ferreira (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 28:1–28:16. <https://doi.org/10.4230/LIPIcs.OPODIS.2018.28>
- [19] Rachid Guerraoui and Ron R. Levy. 2004. Robust Emulations of Shared Memory in a Crash-Recovery Model. In *Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS'04) (ICDCS '04)*. IEEE Computer Society, USA, 400–407. <https://doi.org/10.1109/ICDCS.2004.1281605>
- [20] Maurice P. Herlihy and Jeannette M. Wing. 1990. Linearizability: A Correctness Condition for Concurrent Objects. *ACM Trans. Program. Lang. Syst.* 12, 3 (jul 1990), 463–492. <https://doi.org/10.1145/78969.78972>

- [21] Martin Hyland. 1997. Game Semantics. In *Semantics and Logics of Computation*, Andrew M. Pitts and P.Editors Dybjer (Eds.). Cambridge University Press, Cambridge, UK, 131–184. <https://doi.org/10.1017/CBO9780511526619.005>
- [22] Joseph Izraelevitz, Hammurabi Mendes, and Michael L. Scott. 2016. Linearizability of Persistent Memory Objects Under a Full-System-Crash Failure Model. In *Distributed Computing*, Cyril Gavoille and David Ilcinkas (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 313–327. https://doi.org/10.1007/978-3-662-53426-7_23
- [23] Artem Khyzha, Mike Dodds, Alexey Gotsman, and Matthew Parkinson. 2017. Proving Linearizability Using Partial Orders. In *Programming Languages and Systems: 26th European Symposium on Programming, ESOP 2017* (Uppsala, Sweden). Springer-Verlag, Berlin, Heidelberg, 639–667. https://doi.org/10.1007/978-3-662-54434-1_24
- [24] Artem Khyzha and Ori Lahav. 2022. Abstraction for Crash-Resilient Objects. In *Programming Languages and Systems*, Ilya Sergey (Ed.). Springer International Publishing, Cham, 262–289. https://doi.org/10.1007/978-3-030-99336-8_10
- [25] Jérémie Koenig and Zhong Shao. 2020. Refinement-Based Game Semantics for Certified Abstraction Layers. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science* (Saarbrücken, Germany) (LICS '20). Association for Computing Machinery, New York, NY, USA, 633–647. <https://doi.org/10.1145/3373718.3394799>
- [26] Hongjin Liang and Xinyu Feng. 2013. Modular verification of linearizability with non-fixed linearization points. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Seattle, Washington, USA) (PLDI '13). Association for Computing Machinery, New York, NY, USA, 459–470. <https://doi.org/10.1145/2491956.2462189>
- [27] Nancy Lynch and Frits Vaandrager. 1996. Forward and Backward Simulations. *Inf. Comput.* 128, 1 (jul 1996), 1–25. <https://doi.org/10.1006/inco.1996.0060>
- [28] Paul-André Mellies. 2019. Categorical Combinatorics of Scheduling and Synchronization in Game Semantics. *Proc. ACM Program. Lang.* 3, POPL, Article 23 (jan 2019), 30 pages. <https://doi.org/10.1145/3290336>
- [29] Paul-André Mellies and Noam Zeilberger. 2015. Functors are Type Refinement Systems. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Mumbai, India) (POPL '15). Association for Computing Machinery, New York, NY, USA, 3–16. <https://doi.org/10.1145/2676726.2676970>
- [30] Arthur Oliveira Vale, Paul-André Mellies, Zhong Shao, Jérémie Koenig, and Léo Stefanescu. 2022. Layered and Object-Based Game Semantics. *Proc. ACM Program. Lang.* 6, POPL, Article 42 (jan 2022), 32 pages. <https://doi.org/10.1145/3498703>
- [31] Arthur Oliveira Vale, Zhong Shao, and Yixuan Chen. 2023. A Compositional Theory of Linearizability. *Proc. ACM Program. Lang.* 7, POPL, Article 38 (jan 2023), 32 pages. <https://doi.org/10.1145/3571231>
- [32] Arthur Oliveira Vale, Zhong Shao, and Yixuan Chen. 2024. A Compositional Theory of Linearizability. *J. ACM* 71, 2, Article 14 (apr 2024), 107 pages. <https://doi.org/10.1145/3643668>
- [33] Azalea Raad, John Wickerson, Gil Neiger, and Viktor Vafeiadis. 2019. Persistency Semantics of the Intel-X86 Architecture. *Proc. ACM Program. Lang.* 4, POPL, Article 11 (dec 2019), 31 pages. <https://doi.org/10.1145/3371079>
- [34] Azalea Raad, John Wickerson, and Viktor Vafeiadis. 2019. Weak Persistency Semantics from the Ground up: Formalising the Persistency Semantics of ARMv8 and Transactional Models. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 135 (oct 2019), 27 pages. <https://doi.org/10.1145/3360561>
- [35] Uday S. Reddy. 1993. *A Linear Logic Model of State*. Technical Report. Dept. of Computer Science, UIUC, Urbana, IL.
- [36] Uday S. Reddy. 1996. Global State Considered Unnecessary: An Introduction to Object-Based Semantics. *LISP Symb. Comput.* 9, 1 (1996), 7–76. https://doi.org/10.1007/978-1-4757-3851-3_9
- [37] Kasper Svendsen and Lars Birkedal. 2014. Impredicative Concurrent Abstract Predicates. In *Programming Languages and Systems*, Zhong Shao (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 149–168. https://doi.org/10.1007/978-3-642-54833-8_9
- [38] Yuanhao Wei, Naama Ben-David, Michal Friedman, Guy E. Blelloch, and Erez Petrank. 2022. FLiT: a library for simple and efficient persistent algorithms. In *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Seoul, Republic of Korea) (PPoPP '22). Association for Computing Machinery, New York, NY, USA, 309–321. <https://doi.org/10.1145/3503221.3508436>

SUMMARY OF THE APPENDICES

- A** Provides a more detailed account of the sequential, concurrent and crash-aware game models used in the main paper.
- B** Gives our extended treatment of crash-aware linearizability.
- C** Prepares the ground for strict and durable linearizability by discussing the recrash and decrash operations.
- D** Gives our treatment of strict linearizability, including proofs of locality and an observational refinement property for it.
- E** Gives our extended treatment of durable linearizability.
- F** Defines more carefully the kinds of strategies that serve as denotations of the imperative programs in our program logic section.
- G** Gives the full account of our program logic for durable linearizability, including the full definition of our programming language and operational semantics, all the program logic rules, as well as the proof of soundness.
- H** Serves as a continuation of Appendix **G** by giving a small modification of the program logic there so that it shows crash-aware linearizability instead of durable linearizability.
- I** Collects the detailed specification and proofs in our program logic of the FLiT and File examples discussed in the main paper.
- J** Collects all of the proofs omitted elsewhere.

A A CONCURRENT GAME SEMANTICS WITH CRASHES

In this section we define our game model with crashes. This does require a long exposition, and a few different game models. This proves necessary, as durable linearizability involves crash-aware objects as well as objects without crashes, which live in different models of computation. We start by recalling the definitions of sequential games Seq and concurrent games Conc appearing in Oliveira Vale et al. [31], slightly reformulating them in the process (§A.1). Then, we define our concurrent game model with full-system crashes (§A.2). We finish by discussing the issue of neutral elements and define the copycat strategy crashcopy that will play the role of the neutral element in our compositional model (§A.3), and provide a concrete characterization of saturation with respect to crashcopy .

A.1 Concurrent Games

We now recall the sequential and concurrent games used in Oliveira Vale et al. [31]. The sequential game mode is traditional in the game semantics literature [1, 21] except that our presentation differs in, at least at first, not requiring *O-receptivity*.

The concurrent game model is closely related to the model appearing in Ghica and Murawski [16]. We slightly generalize the model in Oliveira Vale et al. [31]. There, concurrent games are *homogenous* in that all agents play the same game locally, i.e. have access to copies the same operations. The model we define here is *heterogenous* in the sense that it allows different agents to play different games locally.

At this point, we must note that our model of concurrent games is parametrized by a set Υ of agent names. We start by defining a notion of move of a game, and what it means to be a well-formed sequential and concurrent play.

Definition A.1 (Move Sets and Well-Formed Plays). We define the set of polarities for sequential games Pol^{seq} and concurrent games Pol^{conc} respectively as the sets:

$$\text{Pol}^{\text{seq}} := \{O, P\} \qquad \text{Pol}^{\text{conc}} := \sum_{\alpha \in \Upsilon} \text{Pol}^{\text{seq}}$$

We define a *move set* as a pair $(M, \lambda : M \rightarrow \text{Pol})$ of a set of *moves* and a polarity assigning function $\lambda : M \rightarrow \text{Pol}$ to a set of *polarities*. Given such a move set, we write M^{pol} for largest subset of M including only moves $m \in M$ such that $\lambda(m) = \text{pol}$.

In the case of a move set $(M, \lambda : M \rightarrow \text{Pol}^{\text{conc}})$ we also write M^α for the largest subset of M including only moves $\lambda(m) = \alpha:\text{pol}$ for any polarity $\text{pol} \in \text{Pol}^{\text{seq}}$. Note that M^α defines a move set $(M^\alpha, \lambda^\alpha : M \rightarrow \text{Pol}^{\text{seq}})$ by the polarity assignment

$$\lambda^\alpha(m) = \text{pol} \iff \lambda(m) = \alpha:\text{pol}$$

Given a sequence $s \in M^*$ and a subset $S \subseteq \Upsilon$, we denote by $\pi_\alpha(s) \in M^*$ the largest subsequence of s containing only moves in M^α .

Given a move set $(M, \lambda : M \rightarrow \text{Pol}^{\text{seq}})$ we write $\mathbb{P}_M^{\text{seq}}$ for the set of sequences $s \in M^*$ that start with a move labelled by O , and alternate between O and P moves in the sense that

$$s = p \cdot m \cdot m' \cdot t \implies \lambda(m) \neq \lambda(m')$$

We call sequences in $\mathbb{P}_M^{\text{seq}}$ well-formed sequential plays.

Similarly, given a move set $(M, \lambda : M \rightarrow \text{Pol}^{\text{conc}})$ we write $\mathbb{P}_M^{\text{conc}}$ for the set of sequences $s \in M^*$ which are locally sequential in that $\pi_\alpha(s) \in \mathbb{P}_{M^\alpha}^{\text{seq}}$. We call sequences in $\mathbb{P}_M^{\text{conc}}$ well-formed concurrent plays.

We are now ready to define the key notion of a game.

Definition A.2 (Sequential and Concurrent Games). A sequential game $A = ((M_A, \lambda_A), P_A)$ consists of a move set $(M_A, \lambda_A : M_A \rightarrow \text{Pol}^{\text{seq}})$ and a non-empty, prefix-closed, subset $P_A \subseteq \mathbb{P}_A^{\text{seq}}$, where we write $\mathbb{P}_A^{\text{seq}}$ for $\mathbb{P}_{M_A}^{\text{seq}}$.

A concurrent game $A = ((M_A, \lambda_A), P_A)$ consists of a move set $(M_A, \lambda_A : M_A \rightarrow \text{Pol}_\Upsilon^{\text{conc}})$ and a non-empty, prefix-closed subset $P_A^\alpha \subseteq \mathbb{P}_{M_A^\alpha}^{\text{seq}}$ for each $\alpha \in \Upsilon$ verifying that $P_A = \prod_{\alpha \in \Upsilon} P_A^\alpha$.

Given a game A we can recover the local game that α plays, written A^α , as the sequential game $A^\alpha = ((M_A^\alpha, \lambda_A^\alpha), \pi_\alpha(P_A))$. Note that by construction, given $s \in P_A$ it is necessarily the case that $\pi_\alpha(s) \in P_{A^\alpha}$.

Conversely, given an Υ -indexed collection of sequential games $A = (A[\alpha])_{\alpha \in \Upsilon}$ we define the concurrent game $\text{Conc } A$ by the data

$$M_{\text{Conc } A} := \sum_{\alpha \in \Upsilon} M_{A[\alpha]} \quad \lambda_{\text{Conc } A} := \sum_{\alpha \in \Upsilon} \lambda_{A[\alpha]} \quad P_{\text{Conc } A} := \prod_{\alpha \in \Upsilon} \alpha:P_{A[\alpha]}$$

when the context permits, we write A instead $\text{Conc } A$. It is useful to note that given a concurrent game A , we always have $A \cong (A^\alpha)_{\alpha \in \Upsilon}$, where the isomorphism holds up to renaming some moves. So, up-to isomorphism, every concurrent game may be seen as a collection $A = (A[\alpha])_{\alpha \in \Upsilon}$ of sequential games.

An example of a sequential game is the unit game Σ , in which Opponent is able to ask a question q and Proponent may answer the unique answer available to it a . Formally, its moves are $M_\Sigma := \{q, a\}$ where q is an Opponent move, i.e. $\lambda_\Sigma(q) = O$, and a is a Proponent move, i.e. $\lambda_\Sigma(a) = P$, while its set of plays is given by

$$P_\Sigma := \{\epsilon, q, q \cdot a\}$$

We can use the sequential unit game to define an Υ -indexed collection Σ with $\Sigma[\alpha] := \Sigma$. The corresponding game Σ simply allows all agents in Υ to play an instance of Σ locally, so that its plays are sequentially consistent interleavings of plays of Σ labelled by the agent issuing the corresponding moves. For example, for $\alpha \neq \alpha'$ both agents in Υ , the following sequence a play of Σ :

$$\alpha:q \xrightarrow{\quad} \alpha':q \xrightarrow{\quad} \alpha':a \xrightarrow{\quad} \alpha:a$$

where the arrows are merely a visual aid in keeping track of the individual threads of computation.

Definition A.3 (Strategies). For a sequential/concurrent game A , a sequential/concurrent *strategy* σ over A , written $\sigma : A$, is a non-empty, prefix-closed subset $\sigma \subseteq P_A$.

Definition A.4 (Dual Move Set). Given a move set $(M, \lambda : M \rightarrow \text{Pol}^{\text{seq}})$ we define $(M^\perp, \lambda^\perp : M^\perp \rightarrow \text{Pol}^{\text{seq}})$ by

$$M^\perp := M \qquad \lambda^\perp(m) := (\lambda(m))^\perp$$

where $O^\perp := P$ and $P^\perp := O$.

Given a move set $(M, \lambda : M \rightarrow \text{Pol}^{\text{conc}})$ we define $(M^\perp, \lambda^\perp : M^\perp \rightarrow \text{Pol}^{\text{conc}})$ by

$$M^\perp := M \qquad \lambda^\perp(m) := (\lambda(m))^\perp$$

where $(\alpha:\text{pol})^\perp := \alpha:\text{pol}^\perp$ for $\text{pol} \in \text{Pol}^{\text{seq}}$.

Definition A.5 (Tensor and Affine Implication). Given sequential/concurrent games A and B we define the sequential/concurrent games $A \otimes B$ and $A \multimap B$ by the following data

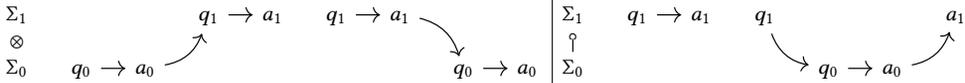
$$M_{A \otimes B} := M_A + M_B \quad \lambda_{A \otimes B} := \lambda_A + \lambda_B \quad P_{A \otimes B} := \{s \in \mathbb{P}_{A \otimes B} \mid s \upharpoonright_{A,-} \in P_A \wedge s \upharpoonright_{-,B} \in P_B\}$$

$$M_{A \multimap B} := M_A^\perp + M_B \quad \lambda_{A \multimap B} := \lambda_A^\perp + \lambda_B \quad P_{A \multimap B} := \{s \in \mathbb{P}_{A \multimap B} \mid s \upharpoonright_{A,-} \in P_A \wedge s \upharpoonright_{-,B} \in P_B\}$$

where $-\upharpoonright_{A,-} : (M_A + M_B)^* \rightarrow M_A$ assigns to a sequence s its largest subsequence involving only elements in M_A , and analogously for $-\upharpoonright_B : (M_A + M_B)^* \rightarrow M_B$.

The plays of $A \otimes B$ are essentially plays of A and B interleaved in a sequentially consistent way, so that $A \otimes B$ corresponds to independent horizontal composition. The game $A \multimap B$ meanwhile corresponds to switching the roles of Opponent and Proponent in A and then taking the tensor with B , which is the traditional way of modelling the affine implication type in game models.

As a matter of illustration, the maximal plays (under prefix ordering) for the sequential games $\Sigma_0 \otimes \Sigma_1$ (the two plays on the left) and $\Sigma_0 \multimap \Sigma_1$ (the two plays on the right) are depicted below. We denote by Σ_0, Σ_1 the two components of these types, both of which are instances of the game Σ . We will similarly add an index to the moves of each component.



Observe that in the game $\Sigma \otimes \Sigma$ Opponent can choose to start in either component, while in the game $\Sigma \multimap \Sigma$ Opponent must start in the target component (Σ_1) due to the flip of polarity in the source component (Σ_0). In $\Sigma \otimes \Sigma$ only Opponent may switch components, while in $\Sigma \multimap \Sigma$ only Proponent may switch components because of alternation (these are typically called the switching conditions of sequential games). Their concurrent variants $\Sigma \otimes \Sigma$ and $\Sigma \multimap \Sigma$ merely allow for any sequentially consistent interleaving of the plays above, labelled by the agents who are performing each move.

Definition A.6. For sequential/concurrent games A, B and C , we define the set of sequential/concurrent interaction plays as

$$\text{int}(A, B, C) := \{s \in (M_A + M_B + M_C)^* \mid s \upharpoonright_{A,B,-} \in P_{A \multimap B} \wedge s \upharpoonright_{-,B,C} \in P_{B \multimap C}\}$$

where $-\upharpoonright_{A,B,-} : (M_A + M_B + M_C)^* \rightarrow (M_A + M_B)^*$ assigns to s the largest subsequence of s involving only events in M_A and M_B , and analogously for the projection $-\upharpoonright_{-,B,C}$.

Given sequential/concurrent strategies $\sigma : A \multimap B$ and $\tau : B \multimap C$ we define their set of interactions by

$$\text{int}(\sigma, \tau) := \{s \in \text{int}(A, B, C) \mid s \upharpoonright_{A,B,-} \in \sigma \wedge s \upharpoonright_{-,B,C} \in \tau\}$$

and their composition, for sequential strategies and concurrent strategies, respectively:

$$\sigma; \tau := \{s \upharpoonright_{A,-,C} \in \mathbb{P}_{A \rightarrow C}^{\text{seq}} \mid s \in \text{int}(\sigma, \tau)\} \quad \sigma; \tau := \{s \upharpoonright_{A,-,C} \in \mathbb{P}_{A \rightarrow C}^{\text{conc}} \mid s \in \text{int}(\sigma, \tau)\}$$

where, $-\upharpoonright_{A,-,C} : (M_A + M_B + M_C)^* \rightarrow (M_A + M_C)^*$ assigns to s the largest subsequence of s containing only moves in M_A and M_C .

PROPOSITION A.7. *Composition of sequential/concurrent strategies is well-defined and associative.*

Prop. A.7, proved by Oliveira Vale et al. [31], means that sequential games, and concurrent games both assemble into semicategories. Recall that a semicategory is essentially a category without the requirements about the neutral element for composition.

Definition A.8. We call **Seq** the semicategory of sequential games and sequential strategies, and **Conc** the semicategory of concurrent games and concurrent strategies.

An important class of games for us, as they will make for the types of our concrete objects and will play a key rule in defining the semantics of imperative code, are games generated by effect signatures. This follows an approach for modeling imperative programs started in Koenig and Shao [25], Oliveira Vale et al. [30]. First, we recall the definition of effect signature and define a notion of concurrent effect signature.

Definition A.9. An effect signature is given by a collection of operations, or effects, $E = (e_i)_{i \in I}$ together with an assignments $\text{par}(-), \text{ar}(-) : E \rightarrow \text{Set}$ of a set of parameters $\text{par}(e)$ and a set of return values $\text{ar}(e)$ for each operation $e \in E$. This is conveniently described by the following notation:

$$E = \{e_i : \text{par}(e_i) \rightarrow \text{ar}(e_i) \mid i \in I\}$$

For a given set of agents Υ we call an Υ -indexed collection of effect signatures $E = (E[\alpha])_{\alpha \in \Upsilon}$ a concurrent effect signature.

The corresponding games associated sequential and concurrent games associated to effect signatures are as follows.

Definition A.10. Given an effect signature E , define the game **Seq**(E) \in **Seq** by the data:

$$M_{\text{Seq}(E)} := \sum_{e \in E} \text{par}(e) + \sum_{e \in E} \text{ar}(e) \quad \lambda_{\text{Seq}(E)} := O + P \quad P_{\text{Seq}(E)} := \downarrow \cup_{e \in E} \cup_{a \in \text{par}(e)} \cup_{v \in \text{ar}(e)} e(a) \cdot v$$

we take the freedom of writing E for **Seq**(E).

Then, given a concurrent effect signature E , its corresponding game **Conc**(E) in **Conc** is the game **Conc** E . We will often write just E for **Conc** E .

The typical play of E looks like the following (on the left), where $e \in E$ is an effect, $a \in \text{par}(e)$ and $v \in \text{ar}(e)$.

$$E : e(a) \longrightarrow v \quad \dagger E : e_1(a_1) \rightarrow v_1 \rightarrow e_2(a_2) \rightarrow v_2 \rightarrow \dots \rightarrow e_n(a_n) \rightarrow v_n$$

Note that E only allows for a single effect of E to be issued. We can lift such a game E to a game $\dagger E$ (read “replay E ”) that allows several effects of E to be invoked in sequence. Its plays, depicted above on the right, consist of sequences of invocations $e_i \in E$ with argument $a_i \in \text{par}(e_i)$ alternating with their responses $v_i \in \text{ar}(e_i)$. We will define the replay modality \dagger – formally later. This informal description will suffice until then. The concurrent variant E , as usual, just allows each agent to play an instance of the corresponding sequential game $E[\alpha]$ in a sequentially consistent fashion, while $\dagger E$ similarly allows each agent to play $\dagger E[\alpha]$ locally.

A.2 Games with Full-System Crashes

We are now ready to define our model of concurrent computation with full-system crashes. We follow the same structure as §A.1.

Definition A.11 (Polarities with Crashes). We define the set of polarities for crash-aware games Pol^{\sharp} as the set $\text{Pol}^{\sharp} := \text{Pol}^{\text{conc}} + \{\sharp\}$.

Note that crash-aware games will effectively have an extra player which is responsible for issuing crash signals, which are represented by moves labelled by \sharp and will be treated differently from O and P moves.

Definition A.12 (Move Set and Well-Formed Plays). We define a crash-aware move set to be a move set $(M, \lambda : M \rightarrow \text{Pol}^{\sharp})$ such that M^{\sharp} is a singleton set.

We write M^{Υ} for the largest subset of M including only moves $m \in M$ such that $\lambda(m) \neq \sharp$, i.e. $M^{\Upsilon} = \cup_{\alpha \in \Upsilon} M^{\alpha}$. Note that M^{Υ} always defines a move set $(M^{\Upsilon}, \lambda^{\Upsilon} : M^{\Upsilon} \rightarrow \text{Pol}^{\text{conc}})$ by the polarity assignment $\lambda^{\Upsilon}(m) := \lambda(m)$.

Note that a sequence $s \in M^*$ is of the form

$$s_1 \cdot \sharp \cdot s_2 \cdot \sharp \cdot \dots \cdot \sharp \cdot s_{n+1}$$

where every $s_i \in (M^{\Upsilon})^*$ and $\sharp \in M_A^{\sharp}$. We define $\|s\|$ to be $n + 1$. We also define the operation $\text{epo}_i(-)$ which assigns to s the sequence $\text{epo}_i(s) := s_i$, called the i -th epoch of s . If $i > \|s\|$ then we take the convention that $\text{epo}_i(s) := \epsilon$.

We say a sequence $s \in M^*$ is a well-formed crash-aware play when, for every $i \in \mathbb{N}$, $\text{epo}_i(s) \in \mathbb{P}_{M^{\Upsilon}}^{\text{conc}}$. We denote by \mathbb{P}_M^{\sharp} the set of all well-formed crash-aware plays over the move set $(M, \lambda : M \rightarrow \text{Pol}^{\sharp})$.

Note that in the definition of the well-formedness of crash-aware plays we implicitly already enforce that crashes affect the entire system, as when a crash happens the entire system resets back to a P -position. This means that after a crash, the first move for all agents is an O move, no matter what was the last move by that agent in the previous epoch. We are now ready to define crash-aware games. In the definition, and for the remainder of this paper, we will make frequent use of the usual Kleene algebra over sequences. We will also denote by \sharp the unique crash event $\sharp \in M^{\sharp}$, for any crash-aware move set, what allows us to ignore the actual name of the crash event. While the assumption that M_A^{\sharp} is a singleton is not necessary, allowing for several crash events makes several of the definitions more involved. Since for our practical purposes one crash event is enough, we opt for this simpler presentation.

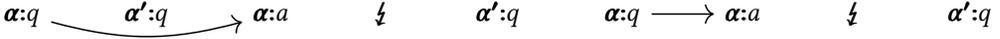
Definition A.13 (Crash-Aware Game). A crash-aware game $A = (M_A, \lambda_A, P_A)$ consists of a move set $(M_A, \lambda_A : M_A \rightarrow \text{Pol}^{\sharp})$ and a non-empty, prefix-closed subset $P_A^{\Upsilon} \subseteq \mathbb{P}_{M_A}^{\text{conc}}$ making

$$P_A = (P_A^{\Upsilon} \cdot \sharp)^* \cdot P_A^{\Upsilon} \subseteq \mathbb{P}_{M_A}^{\sharp}$$

Note that any crash-aware game A defines a concurrent game $A^{\Upsilon} := ((M_A^{\Upsilon}, \lambda_A^{\Upsilon}), P_A^{\Upsilon}) \in \text{Conc}$.

Conversely, given a concurrent game $A = ((M_A, \lambda_A), P_A)$ we can construct a crash-aware game $A^{\sharp} := ((M_A + \{\sharp\}, \lambda_A + \sharp), (P_A \cdot \sharp)^* \cdot P_A)$ where we write \sharp for the constant function $\sharp : \{\sharp\} \rightarrow \{\sharp\}$. This game has every agent $\alpha \in \Upsilon$ playing the concurrent game A . It is useful to observe that given a crash-aware game A , $(A^{\Upsilon})^{\sharp} \cong A$.

So, for example, the crash-aware version of Σ is given by Σ^{\sharp} . Then, an example of play of Σ^{\sharp} is:

Fig. 8. Example of a play of Σ_{\downarrow} .

Note that the main well-formedness constraint about plays of crash-aware games is that in each epoch they play a well-formed concurrent play. i.e. a locally sequential play.

Definition A.14 (Crash-Aware Strategy). A strategy $\sigma : A$ over a game A is a non-empty, prefix-closed, subset $\sigma \subseteq P_A$, which is moreover \downarrow -receptive in that

$$\forall s \in \sigma. s \cdot \downarrow \in P_A \implies s \cdot \downarrow \in \sigma$$

The \downarrow -receptivity property of strategies models the usual assumption that crashes may non-deterministically happen at any point in an execution of a program. Surprisingly, it plays a crucial role in proving the symmetric monoidal structure of crash-aware games.

Observe that in the definition of the tensor product and the affine implication for sequential and concurrent games the disjoint union of move sets plays a crucial role. In order to correctly model the instantaneous and synchronous behavior of crashes, we must treat crash signals differently when computing the disjoint union of move sets. For this, we define a smash product which behaves like the disjoint union for O and P moves, but that merges crash signals together. It will also be necessary to redefine projections to take this merger into account.

Definition A.15 (Smash Product). Given move sets (M_A, λ_A) and (M_B, λ_B) we define their smash product $(M_A +_{\downarrow} M_B, \lambda_A +_{\downarrow} \lambda_B)$ by

$$M_A +_{\downarrow} M_B := M_A^{\downarrow} + M_B^{\downarrow} + \downarrow \quad \text{and} \quad \lambda_A +_{\downarrow} \lambda_B := \lambda_A + \lambda_B + \downarrow$$

where \downarrow stands for the constant function to $\downarrow \in \text{Pol}^{\downarrow}$.

Given $s \in \mathbb{P}_{M_A +_{\downarrow} M_B}$ we define $s \upharpoonright_{A,-} \in \mathbb{P}_{M_A}$ and $s \upharpoonright_{-,B} \in \mathbb{P}_{M_B}$ to be the projections to the corresponding components of $M_A +_{\downarrow} M_B$. What this means is that $-\upharpoonright_{M_A,-}$ and $-\upharpoonright_{-,M_B}$ are respectively generated by the maps in **Set** below using the universal property of the free monoids M_A^* and M_B^* respectively:

$$-\upharpoonright_{M_A,-} := \text{id}_{M_A^{\downarrow}} + \epsilon + \text{id}_1 : M_A +_{\downarrow} M_B \rightarrow M_A^* \quad -\upharpoonright_{-,M_B} := \epsilon + \text{id}_{M_B^{\downarrow}} + \text{id}_1 : M_A +_{\downarrow} M_B \rightarrow M_B^*$$

where ϵ is the constant function to the empty sequence, and id_S is the identity for the set S .

Definition A.16 (Dual Move Set). Given a move set (M, λ) we define the moveset $(M^{\perp}, \lambda^{\perp})$ by

$$M^{\perp} := M \quad \lambda^{\perp}(m) := \lambda(m)^{\perp}$$

where $(\alpha:\text{pol})^{\perp} := \alpha:\text{pol}^{\perp}$ for $\text{pol} \in \text{Pol}^{\text{seq}}$, and $\downarrow^{\perp} := \downarrow$.

In the context of games A and B , as opposed to move sets, we write $-\upharpoonright_{A,-}$ and $-\upharpoonright_{-,B}$ for $-\upharpoonright_{M_A,-}$ and $-\upharpoonright_{-,M_B}$ respectively. We will also write $-\upharpoonright_{A,B,-}$ and $-\upharpoonright_{-,B,C}$ for, respectively,

$$-\upharpoonright_{M_A +_{\downarrow} M_B, -} : (M_A +_{\downarrow} M_B +_{\downarrow} M_C)^* \rightarrow (M_A +_{\downarrow} M_B)^* \quad -\upharpoonright_{-, M_B +_{\downarrow} M_C} : (M_A +_{\downarrow} M_B +_{\downarrow} M_C)^* \rightarrow (M_B +_{\downarrow} M_C)^*$$

We also take the opportunity to define a projection

$$-\upharpoonright_{A,-,C} : (M_A +_{\downarrow} M_B +_{\downarrow} M_C)^* \rightarrow (M_A +_{\downarrow} M_C)^*$$

as the monoid homomorphism associated by the universal property of the free monoid $(M_A +_{\downarrow} M_C)^*$ to the mapping $p_{A,-,C} : M_A +_{\downarrow} M_B +_{\downarrow} M_C \rightarrow (M_A +_{\downarrow} M_C)^*$ in **Set**:

$$p_{A,-,C}(m) = \begin{cases} m, & m \in M_A^{\downarrow} + M_C^{\downarrow} \\ \epsilon, & m \in M_B^{\downarrow} \\ \downarrow, & (\lambda_A +_{\downarrow} \lambda_B +_{\downarrow} \lambda_C)(m) = \downarrow \end{cases}$$

Definition A.17. Fix games A and B . We define the games $A \otimes B$ and $A \multimap B$ by the following data

$$\begin{aligned} M_{A \otimes B} &:= M_A +_{\downarrow} M_B & \lambda_{A \otimes B} &:= \lambda_A +_{\downarrow} \lambda_B & P_{A \otimes B} &:= \{s \in \mathbb{P}_{M_A +_{\downarrow} M_B} \mid s \upharpoonright_{A,-} \in P_A \wedge s \upharpoonright_{-,B} \in P_B\} \\ M_{A \multimap B} &:= M_A^{\perp} +_{\downarrow} M_B & \lambda_{A \multimap B} &:= \lambda_A^{\perp} +_{\downarrow} \lambda_B & P_{A \multimap B} &:= \{s \in \mathbb{P}_{M_A^{\perp} +_{\downarrow} M_B} \mid s \upharpoonright_{A,-} \in P_A \wedge s \upharpoonright_{-,B} \in P_B\} \end{aligned}$$

It is in this change from a disjoint union to the smash product, and in the corresponding modification to projections that lies some of the assumptions about the behavior of crashes. Consider the following play s of $\Sigma^{\downarrow} \multimap \Sigma^{\downarrow}$ (on the left):

$$\begin{array}{c} s = \begin{array}{ccc} \alpha:q & & \alpha':q \\ \searrow & \longrightarrow & \downarrow \\ & \alpha:q & \alpha:a \\ & \longrightarrow & \downarrow \\ & & \alpha':q \end{array} \quad \left| \begin{array}{c} \alpha':q \\ \downarrow \\ \alpha':q \end{array} \right. \end{array} \quad \left| \begin{array}{l} s \upharpoonright_{-, \Sigma^{\downarrow}} = \alpha:q \cdot \alpha':q \cdot \downarrow \cdot \alpha':q \\ s \upharpoonright_{\Sigma^{\downarrow}, -} = \alpha:q \cdot \alpha:a \cdot \downarrow \cdot \alpha':q \end{array} \right.$$

Note that the crash signal synchronize across the the source and target components of the play. This simultaneously models that the crashes are *synchronous* across components (they happen in all components at once) and that they are *instantaneous* (it takes negligible time for the crash to propagate to other components). On the right, above, we see the projections of s to the source and target components. Importantly, the crash event is retained in both projections, so that the crash event \downarrow of $\Sigma^{\downarrow} \multimap \Sigma^{\downarrow}$ effectively belongs to both components.

Definition A.18. We define the set of interaction plays as

$$\text{int}(A, B, C) := \{s \in (M_A +_{\downarrow} M_B +_{\downarrow} M_C)^* \mid s \upharpoonright_{A,B,-} \in P_{A \multimap B} \wedge s \upharpoonright_{-,B,C} \in P_{B \multimap C}\}$$

Given strategies $\sigma : A \multimap B$ and $\tau : B \multimap C$ we define their set of interactions by

$$\text{int}(\sigma, \tau) := \{s \in \text{int}(A, B, C) \mid s \upharpoonright_{A,B,-} \in \sigma \wedge s \upharpoonright_{-,B,C} \in \tau\}$$

and their composition

$$\sigma; \tau := \{s \upharpoonright_{A,-,C} \in \mathbb{P}_{A \multimap C}^{\downarrow} \mid s \in \text{int}(\sigma, \tau)\}$$

PROPOSITION A.19. *Composition of crash-aware strategies is well-defined and associative.*

Definition A.20. We denote by **Crash** the semicategory of crash-aware games, with crash-aware strategies $\sigma : A \multimap B$ as morphisms between games A and B , and composition given by $;$, $-$.

We take the opportunity to define the crash-aware version of E as E^{\downarrow} . We similarly define the game $\dagger E^{\downarrow}$ as $(\dagger E)^{\downarrow}$.

A.3 The Copycat Strategies and Saturation

None of **Seq**, **Conc** or **Crash** assemble into categories for the corresponding composition operations $;$, $-$ do not have a neutral element. That is, to say, there is no choice of strategies $\text{id}_A : A \multimap A$ for which $\text{id}_A; \sigma; \text{id}_B = \sigma$ for every $\sigma : A \multimap B$. On the other hand, there are clear candidates for such a neutral element, which are called the copycat strategies.

The sequential copycat strategy $\text{seqcopy}_A : A \multimap A$ intuitively replicates O -moves in the target component as the same move in the source component, and P -moves in the source component to the same move in the target component. Formally, it is defined by:

$$\text{seqcopy}_A := \{s \in P_{A \multimap A} \mid \forall p \sqsubseteq_{\text{even}} s. p \upharpoonright_{A_1} = p \upharpoonright_{A_2}\}$$

where we write \sqsubseteq for the prefix relation and $\sqsubseteq_{\text{even}}$ for the even-length prefix relation. For Σ the maximal play of the copycat strategy seqcopy_{Σ} is the play below (on the left):

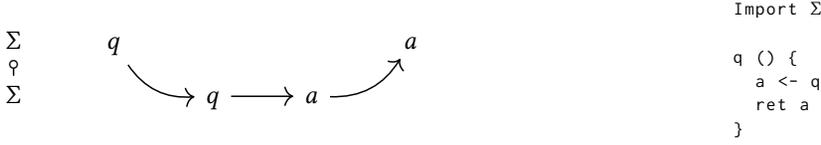


Fig. 9. Maximal play of seqcopy_Σ (left) and corresponding pseudocode (right)

This play corresponds to complete execution of the code we display on the right, which implements Σ by importing another instance of Σ .

The concurrent copycat then $\text{ccopy}_A : A \multimap A$ merely has every agent $\alpha \in \Upsilon$ play the sequential copycat for its corresponding game $\text{seqcopy}_{A^\alpha} : A^\alpha \multimap A^\alpha$:

$$\text{ccopy}_A := \{s \in P_{A \multimap A} \mid \pi_\alpha(s) \in \text{seqcopy}_{A^\alpha}\}$$

Finally, the crash-aware copycat $\text{crashcopy}_A : A \multimap A$ just plays $\text{ccopy}_{A^\Upsilon}$ within each epoch:

$$\text{crashcopy}_A := (\text{ccopy}_{A^\Upsilon} \cdot \downarrow)^* \cdot \text{ccopy}_{A^\Upsilon}$$

It turns out that in all cases the corresponding copycat strategy can lead to emergent behavior after composition, which prevents the models from being compositional. Concretely, writing copy generically for any of the copycat strategies, it can happen that for some $\sigma : A \multimap B$, $\sigma \subset \text{copy}_A; \sigma; \text{copy}_B$ strictly. This issue is explained extensively in Oliveira Vale et al. [31] in the context of concurrent games, and the situation is the same for crash-aware games, so we refer the reader there for a more detailed account of the issue.

The solution they present which turns out to be deeply related to linearizability, is to note that the copycat strategy is *idempotent*, in that for all games A in the corresponding models

$$\text{copy}_A; \text{copy}_A = \text{copy}_A$$

This essentially means that the copy_A at least behaves like a neutral element for itself. With that fact at hand, we define a class of strategies that behave well when composed with the copycat.

Definition A.21. We say a strategy $\sigma : A \multimap B$ is saturated with respect to the copycat strategy copy when

$$\text{copy}_A; \sigma; \text{copy}_B = \sigma$$

It is not hard to see that the fact that the copycat strategy is idempotent ensures that composing saturated strategies yields a saturated strategy, and that the copycat does behave like a neutral element for saturated strategies. This means that we can promote the semicategories we have defined to categories by restricting attention to these saturated strategies.

Definition A.22. We define the categories **Seq**, **Conc** and **Crash** as the restrictions, respectively, of the semicategories **Seq**, **Conc** and **Crash** to strategies saturated with respect to the corresponding copycat strategies seqcopy , ccopy and crashcopy .

It is folklore that, concretely, saturation for sequential strategies is equivalent to O -receptivity. That is, a strategy is saturated if and only if it accepts O -moves whenever the environment is allowed to make them. For concurrent games the story for saturation is more complicated, and corresponds to, beyond O -receptivity, strategies that are insensitive to certain delays, which might be caused, for instance, if an agent is preempted. This is typically formalized using a rewrite system $\sim \rightsquigarrow -$, which figures prominently in the concrete formulation of linearizability appearing in Oliveira Vale et al. [31], so naturally it will also play a key role in our own work on linearizability, and we take the opportunity to define it now.

Definition A.23. Let $A = (M_A, P_A)$ be a concurrent game. We define a string rewrite system $(P_A, \rightsquigarrow_A)$ with local rewrite rules:

- $\forall m, m' \in M_A. \forall \alpha, \alpha' \in \Upsilon. \forall X \in \{O, P\}. \alpha \neq \alpha' \wedge \lambda_A(m) = \alpha : X \wedge \lambda_A(m') = \alpha' : X \implies m \cdot m' \rightsquigarrow_A m' \cdot m$
- $\forall o, p \in M_A. \forall \alpha, \alpha' \in \Upsilon. \alpha \neq \alpha' \wedge \lambda_A(o) = \alpha : O \wedge \lambda_A(p) = \alpha' : P \implies o \cdot p \rightsquigarrow_A p \cdot o$

For crash-aware strategies, the concrete characterization is only slightly more involved. We do not cover it here for the sake of space. We will soon see an equivalent characterization in terms of a linearizability criterion, which will be sufficient for our purposes.

A.4 Refinement and Horizontal Composition

Before proceeding, we briefly address refinement and horizontal composition. We take as our notion of refinement behavior containment, $\sigma \subseteq \tau$, with joins given by set union. This makes all of the models we have discussed so far into enriched (semi)categories over join semi-lattices. Specifically, this means that

PROPOSITION A.24. *Strategy composition $-; -$ is monotonic and join-preserving.*

Now, for horizontal composition, recall that we have already defined a game $A \otimes B \in \mathbf{Crash}$ given games A and B in \mathbf{Crash} . The tensor defines a semifunctor as follows, where $\sigma : A \multimap B$ and $\tau : A' \multimap B'$:

$$\sigma \otimes \tau := \{s \in P_{A \otimes A' \multimap B \otimes B'} \mid s \upharpoonright_{A \multimap B} \in \sigma \wedge s \upharpoonright_{A' \multimap B'} \in \tau\}$$

Note that these projections are defined just like in the usual concurrent case except for its behaviour on crashes \downarrow which is given by:

$$\downarrow \upharpoonright_{A \multimap B} = \downarrow \qquad \downarrow \upharpoonright_{A' \multimap B'} = \downarrow$$

Moreover, the game $\mathbf{1}$ is given by the following data

$$M_{\mathbf{1}} = \{\downarrow\} \qquad \lambda_{\mathbf{1}}(\downarrow) = \downarrow \qquad P_{\mathbf{1}} = \downarrow^*$$

These definitions permit us to prove that

PROPOSITION A.25. *(Crash, $- \otimes -$, $\mathbf{1}$) defines an enriched symmetric monoidal category.*

This means, in particular, that $- \otimes -$ defines a monotonic and join-preserving functor, so that horizontal composition behaves well with respect to both vertical composition and refinement.

B CRASH-AWARE LINEARIZABILITY

Compositional linearizability provides an account of linearizability based on an operation $K_{\mathbf{Conc}}$ – converting strategies in \mathbf{Conc} to strategies in \mathbf{Conc} . This operation, comes with the abstract construction defining \mathbf{Conc} , so that, as \mathbf{Crash} follows the same construction, there is a counterpart K_{\downarrow} – in \mathbf{Crash} . In this section, we give a concrete characterization of the notion of linearizability associated to K_{\downarrow} –, which is closely related to strict linearizability. Then, we present the equivalence with observational refinement and the locality property for it.

B.1 Crash-Aware Linearizability

We start by defining the operation $K_{\downarrow} - : \mathbf{Crash} \rightarrow \mathbf{Crash}$ by the formula, for $\tau : A \multimap B \in \mathbf{Crash}$

$$K_{\downarrow} \tau := \text{crashcopy}_A; \sigma; \text{crashcopy}_B$$

Intuitively, this operation assigns to σ the smallest saturated strategy containing σ . $K_{\downarrow} -$ has the important property that it is an oplax semifunctor, i.e.

PROPOSITION B.1. For all $\sigma : A \multimap B$ and $\tau : B \multimap C$ in Crash, $K_{\downarrow}(\sigma; \tau) \subseteq K_{\downarrow} \sigma; K_{\downarrow} \tau$.

The framework of compositional linearizability proposes that the native notion of linearizability for crash-aware objects should be equivalent to the refinement $v' \subseteq K_{\downarrow} v$. So for the remainder of this section, our goal is to concretely characterize this refinement in terms of plays. For this, we find it useful to recall the concrete formulation of compositional linearizability. This notion of linearizability is a slight generalization of interval-sequential linearizability, and in particular does not require that the linearized specification be atomic nor that all pending operations be removed.

Definition B.2. For $A \in \mathbf{Conc}$, a play $s \in P_A$ is linearizable to a play $t \in P_A$ when there exists a sequence of P -moves $s_P \in (M_A^P)^*$ and a sequence of O -moves $s_O \in (M_A^O)^*$ such that $s \cdot s_P \rightsquigarrow_A t \cdot s_O$. We write $s \rightsquigarrow t$ when s is linearizable to t . We say $s \in P_A$ is linearizable with respect to a strategy $\nu : A$, written $s \rightsquigarrow \nu$ when there exists $t \in \nu$ such that $s \rightsquigarrow t$. We say a strategy $\nu' : A$ is linearizable with respect to a strategy $\nu : A$, written $\nu' \rightsquigarrow \nu$, when for every play $s \in \nu'$, $s \rightsquigarrow \nu$.

We use linearizability in Conc to define crash-aware linearizability. It is straight-forward: s crash-aware linearizes to t when their corresponding epochs linearize to each other.

Definition B.3. For $A \in \mathbf{Crash}$, we say a play $s \in P_A$ crash-aware linearizes to a play $t \in P_A$ when

$$\|s\| = \|t\| \quad \text{and} \quad \forall i \leq \|s\|. \text{epo}_i(s) \rightsquigarrow \text{epo}_i(t)$$

and write $s \rightsquigarrow_{\downarrow} t$ when this holds, and extend the notation as we did for linearizability (see Def. B.2).

We now discuss a few examples of crash-aware linearizability. A first example is volatile objects. For this, we find it useful to define a functor $\text{Vol } -$, defined by $\text{Vol } A := A^{\downarrow}$ on games. Meanwhile, given a strategy $\sigma : A \multimap B \in \mathbf{Conc}$ we define the strategy $\text{vol}(\sigma) \in \mathbf{Crash}$ as $\text{vol}(\sigma) := (\sigma \cdot \downarrow)^* \cdot \sigma$. The functor is named $\text{vol}(-)$ because we use it to define volatile objects, as given a strategy $\nu_E : \dagger E \in \mathbf{Conc}$, $\text{vol}(\nu_E) : \dagger E^{\downarrow}$ describes the object that behaves as ν_E within each epoch, and in particular resets the object to its initial state on a crash. It is easy to see that:

PROPOSITION B.4. For $\nu' : A \in \mathbf{Conc}$ and $\nu : A \in \mathbf{Conc}$, if $\nu' \rightsquigarrow \nu$ then $\text{vol}(\nu') \rightsquigarrow_{\downarrow} \text{vol}(\nu)$.

The main result of this section is the following characterization of $K_{\downarrow} -$.

PROPOSITION B.5.

$$K_{\downarrow} \tau = \{s \in P_{A \multimap B} \mid s \text{ is crash-aware linearizable with respect to } \tau\}$$

COROLLARY B.6. $\nu' : A$ is crash-aware linearizable with respect to $\nu : A$ if and only if $\nu' \subseteq K_{\downarrow} \nu$.

B.2 Observational Refinement and Locality

We use the general result in Oliveira Vale et al. [31] to obtain locality and observational refinement. The requirements to obtain these properties are the following.

LEMMA B.7.

- For any $\sigma : 1 \multimap A \in \mathbf{Crash}$ it holds that $\text{crashcopy}_1; \sigma = \sigma$.
- For $\sigma, \tau : A \multimap B$ and $\sigma', \tau' : A' \multimap B'$ we have $\sigma \otimes \sigma' \subseteq \tau \otimes \tau' \implies \sigma \subseteq \sigma' \wedge \tau \subseteq \tau'$

This gives as corollaries locality and observational refinement, which we write explicitly now. The proof of these results, under the conditions of our construction and Lemma B.7.

COROLLARY B.8 (OBSERVATIONAL REFINEMENT). $\nu'_A : A \in \mathbf{Crash}$ is crash-aware linearizable w.r.t $\nu_A : A \in \mathbf{Crash}$ if and only if for all $\sigma : A \multimap B$, $\nu'_A; \sigma \subseteq \nu_A; \sigma$

COROLLARY B.9 (LOCALITY). For $\nu'_A : A, \nu'_B : B \in \mathbf{Crash}$ and $\nu_A : A, \nu_B : B \in \mathbf{Crash}$:

$$\nu'_A \rightsquigarrow_{\downarrow} \nu_A \text{ and } \nu'_B \rightsquigarrow_{\downarrow} \nu_B \text{ if and only if } \nu'_A \otimes \nu'_B \rightsquigarrow_{\downarrow} \nu_A \otimes \nu_B$$

C CRASH ABSTRACTION

Many specification methodologies for crash-aware objects, including durable linearizability, use specifications without crashes. In our framework, this means that while the concrete crash-aware object v' lives in Vol , the abstract specification v lives in Conc . In this section we develop conversions between **Crash** and **Conc** that serve as a building block for strict and durable linearizability.

The main difficulty in removing crashes from a play s with crashes is that the removal may generate traces not satisfying sequential consistency. This happens when the same agent has a pending invocation in one epoch and also moves in a later epoch. So, in the definition of the operation $-^b$ (read *de-crash*), the projections $\pi_\Upsilon(s)$ are required to be well-formed plays.

Definition C.1. Given a game $A = (M_A, \lambda_A, P_A) \in \text{Crash}$ we define the game $A^b \in \text{Conc}$, by:

$$M_{A^b} := M_A^\Upsilon \quad \lambda_{A^b}(m) := \lambda_A(m) \quad P_{A^b} := \{\pi_\Upsilon(s) \in \mathbb{P}_{A^b}^{\text{conc}} \mid s \in \|\alpha \in \Upsilon (P_A^\alpha)^*\}$$

Given a strategy $\sigma : A \in \text{Crash}$ we similarly define $\sigma^b : A^b \in \text{Conc}$ by

$$\sigma^b := \{\pi_\Upsilon(s) \in \mathbb{P}_{A^b}^{\text{conc}} \mid s \in \sigma\}$$

A result we note in passing is that $(A \multimap B)^b \cong A^b \multimap B^b$. This permits us to show that $-^b$ defines an oplax semifunctor $-^b : \text{Crash} \rightarrow \text{Conc}$, that is:

PROPOSITION C.2. For all $\sigma : A \multimap B$ and $\tau : B \multimap C$ in **Crash**, $(\sigma; \tau)^b \subseteq \sigma^b; \tau^b$.

In addition, for all $A \in \text{Crash}$, $\text{crashcopy}_A^b = \text{ccopy}_A$, and $-^b$ is monotonic and join-preserving.

It is also useful to provide a reverse operation $-^\sharp$, read *re-crash*, that lifts, in a persistent way, a strategy $\sigma : A^b \multimap B^b$ into a strategy $\sigma^\sharp : A \multimap B$, a situation we depict diagrammatically below.

$$\begin{array}{ccc} A & \overset{\sigma^\sharp}{\multimap} & B \\ \downarrow & & \downarrow \\ A^b & \overset{\sigma}{\multimap} & B^b \end{array}$$

So, suppose given $A, B \in \text{Crash}$ and $\sigma : A^b \multimap B^b \in \text{Conc}$. Then, the strategy $\sigma^\sharp : A \multimap B \in \text{Crash}$ is given by:

$$\sigma^\sharp := \{s \in \mathbb{P}_A^\sharp \mid \pi_\Upsilon(s) \in \sigma\}$$

Re-crash also behaves *like* an enriched semifunctor.

PROPOSITION C.3. For $\sigma : A^b \multimap B^b, \tau : B^b \multimap C^b \in \text{Conc}$, $\sigma^\sharp; \tau^\sharp \subseteq (\sigma; \tau)^\sharp$.

In addition, for all $A \in \text{Crash}$, $\text{ccopy}_{A^b}^\sharp \subseteq \text{crashcopy}_A$, and $-^\sharp$ is monotonic and join-preserving.

D STRICT LINEARIZABILITY

Strict linearizability [2], an important linearizability criterion in the presence of crashes often used to specify objects with robust recovery routines, postulates that a pending operation must linearize within the same epoch it was issued. It was originally formulated in a system with individual crashes instead of full-system crashes, so we modify it for our setting. We do not assume atomicity or that all pending operations are removed.

Similarly to how Oliveira Vale et al. [31] characterizes linearizability by lifting a non-saturated strategy to a saturated strategy, we formalize strict linearizability by lifting a strategy without crashes into a strategy with crashes. Naively, one might think it is enough to use the lift $v^\sharp : A \in \text{Crash}$. Unfortunately, this does not make a crash-aware object, mainly because it does not satisfy O -receptivity anymore. Specifically, v^\sharp never has plays such as

$$\alpha : q \cdot \frac{1}{2} \cdot \alpha : q \cdot \alpha : a$$

because $\alpha : q \cdot \alpha : q \cdot \alpha : a$ is not well-formed as a play of A^b . We can fix this issue by saturating the resulting strategy using $K_{\frac{1}{2}} -$.

Definition D.1. Given games $A, B \in \underline{\mathbf{Crash}}$, we define the strict lift $\text{str}(\sigma) : A \multimap B$ of a strategy $\sigma : A^b \multimap B^b \in \underline{\mathbf{Conc}}$ as the strategy

$$\text{str}(\sigma) := K_{\frac{1}{2}} \sigma^{\sharp}$$

We now define our variation of strict linearizability, which differs only in that it generalizes the original strict-linearizability by not requiring the linearized specification to be atomic and complete, and specializes it to full-system crashes.

Definition D.2. We say $v' : A \in \mathbf{Crash}$ is strictly linearizable to $v : A^b \in \underline{\mathbf{Conc}}$ when $v' \subseteq \text{str}(v)$.

It is not hard to see that when v is an atomic strategy this is equivalent to strict linearizability. Indeed, the application of $K_{\frac{1}{2}}$ – within $\text{str}(-)$ corresponds to constructing a strict completion [3], and $-^{\sharp}$ to removing the crashes. Finally, by first noting that

PROPOSITION D.3. For all $\sigma : A^b \multimap B^b, \tau : B^b \multimap C^b \in \underline{\mathbf{Conc}}$, $\text{str}(\sigma); \text{str}(\tau) \subseteq \text{str}(\sigma; \tau)$.

we are able to prove the following refinement property for strict linearizability.

PROPOSITION D.4. Suppose $v'_A : A$ is strictly linearizable to $v_A : A^b$ and that $\sigma : A^b \multimap B^b$ implements an object linearizable to $v_B : B^b$ using v_A , i.e.

$$v_A; \sigma \subseteq v_B$$

Then, $\text{str}(\sigma)$ implements an object strictly linearizable to $\text{str}(v_B)$ using v'_A , i.e.

$$v'_A; \text{str}(\sigma) \subseteq \text{str}(v_B)$$

The reverse direction, unfortunately does not hold as $\text{str}(\text{ccopy}_{A^b}) \neq \text{crashcopy}_{A^b}$. By similar reasoning as the locality for crash-aware linearizability we also obtain that

PROPOSITION D.5 (LOCALITY). For $v'_A : A, v'_B : B \in \mathbf{Crash}$ and $v_A : A, v_B : B \in \underline{\mathbf{Conc}}$:

$$v'_A \subseteq \text{str}(v_A) \text{ and } v'_B \subseteq \text{str}(v_B) \text{ if and only if } v'_A \otimes v'_B \subseteq \text{str}(v_A \otimes v_B)$$

E DURABLE LINEARIZABILITY

A frequently used linearizability criterion for specifying persistent objects is durable linearizability [22], which appears as an important development in linearizability with crashes [3]. Durable linearizability makes a core assumption, which we call the *durability assumption*: that the agents appearing in each epoch are disjoint across epochs. This means that there is no agent re-use between epochs. The assumption makes the definition of durable linearizability significantly simpler than previous criteria, including strict linearizability. Moreover, under the durability assumption, persistent atomicity and recoverable linearizability are equivalent. This assumption, however, is quite intrusive and must be enforced throughout, requiring us to define a new model Dur . Interestingly, this can be smoothly done by a different choice of copycat for $\underline{\mathbf{Crash}}$. We finish by showing locality and observational refinement for durable linearizability.

E.1 Durable Linearizability

We start by formalizing the durability assumption, and from now on assume Υ is countably infinite.

Definition E.1. For a move set $(M_A, \lambda_A : M_A \rightarrow \text{Pol}^{\frac{1}{2}})$ we say a play $s \in \mathbb{P}_A$ is *durable* when the set of agents appearing in each epoch of s are pairwise disjoint across epochs, i.e.

$$\forall i. \forall j. i \neq j \implies \Upsilon(\text{epo}_i(s)) \cap \Upsilon(\text{epo}_j(s)) = \emptyset$$

and write $\mathbb{P}_A^{\text{dur}}$ for the set of all durable plays over M_A . We write P_A^{dur} for $P_A \cap \mathbb{P}_A^{\text{dur}}$.

We say a strategy is *durable* when all of its plays are durable.

Durable plays s have the important property that $\pi_Y(s) \in P_A$ as

$$\pi_Y(s) = \text{epo}_1(s) \cdot \text{epo}_2(s) \cdot \dots \cdot \text{epo}_n(s) \cdot \dots$$

writing $\text{ops}(s)$ for the right-hand side above we obtain that when σ is durable $\sigma^b = \text{ops}(\sigma)$. An important result is that durable strategies compose:

PROPOSITION E.2. *If $\sigma : A \multimap B$ and $\tau : B \multimap C$ are durable strategies then $\sigma; \tau$ is a durable strategy.*

This means that the restriction of **Crash** to durable strategies defines a semicategory, which we call **Dur**. This motivates defining a durable copycat strategy.

Definition E.3. For $A \in \text{Crash}$, the strategy $\text{durcopy}_A : A \multimap A$ is the strategy:

$$\text{durcopy}_A := \text{crashcopy}_A \cap \mathbb{P}_{A \multimap A}^{\text{dur}}$$

PROPOSITION E.4. *durcopy is idempotent.*

This let's us define a model for durable objects and their implementations.

Definition E.5. We define the category **Dur** as the restriction of **Crash** to strategies saturated with respect to durcopy.

The construction, like for **Conc** and **Crash**, comes with its own operation $K_{\text{dur}} : \text{Dur} \rightarrow \text{Dur}$ defined as $K_{\text{dur}} \sigma := \text{durcopy}_A; \sigma; \text{durcopy}_B$, as expected. Its associated notion of linearizability $\nu' \subseteq K_{\text{dur}} \nu$ is the same as crash-aware linearizability restricted to durable strategies. An important fact that we mention in passing is that restriction to durable plays $- \cap \mathbb{P}_{A \multimap B}$ defines a semifunctor from **Crash** to **Dur**. Composing this semifunctor with $-^\#$ therefore preserves its functoriality properties. For simplicity, we denote this composition simply as $-^\#$, as the context should make it clear when the durable assumption is in place. We are finally ready to define durable linearizability.

Definition E.6. We say a play $s \in P_A^{\text{dur}}$ is durably linearizable to a play $t \in P_{A^b}$, written $s \stackrel{\text{dur}}{\rightsquigarrow} t$, when $\text{ops}(s) \rightsquigarrow t$. We extend the notation to strategies as we did for linearizability (see Def. B.2).

Our definition differs from the traditional one only in that: we do not assume the linearized trace is atomic, and we do require that all pending invocations be removed in the linearization.

Now, for our refinement-based formulation, we define a durable lift $\text{dur}(-)$, which assigns to a strategy $\nu : A^b \in \text{Conc}$ the strategy $\text{dur}(\nu) : A \in \text{Dur}$ defined by

$$\text{dur}(\nu) : A \in \text{Crash} := (K_{\text{Conc}} \nu)^\# \cap P_A^{\text{dur}}$$

And, indeed, $\text{dur}(-)$ does provide an appropriate lifting operation for durable linearizability.

PROPOSITION E.7. *$\nu' : A \in \text{Crash}$ is durable linearizable to $\nu : A^b \in \text{Conc}$ if and only if $\nu' \subseteq \text{dur}(\nu)$.*

E.2 Observational Refinement and Locality

E.2.1 Observational Refinement. We now show an observational refinement property for durable linearizability. For this, we first note that $\text{dur}(-)$ behaves *like* a lax semifunctor, that is:

PROPOSITION E.8. *For all $\sigma : A^b \multimap B^b, \tau : B^b \multimap C^b \in \text{Conc}$, $\text{dur}(\sigma); \text{dur}(\tau) \subseteq \text{dur}(\sigma; \tau)$.*

It also satisfies the following property, which $\text{str}(-)$ does not satisfy:

PROPOSITION E.9. *For all $A \in \text{Dur}$, $\text{dur}(\text{ccopy}_{A^b}) = \text{durcopy}_A$*

Each of these two results play an important role in each of the two directions of the following equivalence with observational refinement.

PROPOSITION E.10. Let $A, B \in \mathbf{Crash}$. Then $v'_A : A$ is durably linearizable to $v_A : A^b$ if and only if whenever $\sigma : A^b \multimap B^b \in \mathbf{Conc}$ implements a concurrent object linearizable to v_B using v_A , then $\text{dur}(\sigma) : A \multimap B$ implements an object durably linearizable to v_B using v'_A .

E.2.2 *Locality*. For locality, we start by defining a tensor product.

Definition E.11. For strategies $\sigma : A, \tau : B \in \mathbf{Dur}$ we define their tensor product $\sigma \boxtimes \tau : A \boxtimes B$ as

$$A \boxtimes B := A \otimes B \qquad \sigma \boxtimes \tau := (\sigma \otimes \tau) \cap P_{A \otimes B}^{\text{dur}}$$

PROPOSITION E.12. $(\mathbf{Dur}, - \boxtimes -, \mathbf{1})$ defines a symmetric monoidal category.

With that established, establishing locality follows the same structure as for \mathbf{Conc} and \mathbf{Crash} .

PROPOSITION E.13. For any durable $\sigma : A, \tau : B \in \mathbf{Conc}$, $\text{dur}(\sigma \otimes \tau) = \text{dur}(\sigma) \boxtimes \text{dur}(\tau)$.

PROPOSITION E.14. For $\sigma, \tau : A \multimap B \in \mathbf{Dur}$ and $\sigma', \tau' : A' \multimap B' \in \mathbf{Dur}$ we have

$$\sigma \boxtimes \sigma' \subseteq \tau \boxtimes \tau' \implies \sigma \subseteq \sigma' \wedge \tau \subseteq \tau'$$

COROLLARY E.15 (LOCALITY). For $v'_A : A, v'_B : B \in \mathbf{Dur}$ and $v_A : A, v_B : B \in \mathbf{Conc}$:

$$v'_A \stackrel{\text{dur}}{\sim} v_A \text{ and } v'_B \stackrel{\text{dur}}{\sim} v_B \text{ if and only if } v'_A \boxtimes v'_B \stackrel{\text{dur}}{\sim} v_A \otimes v_B$$

E.3 FLiT Correctness Theorem

For the FLiT correctness theorem, we must assume that v'_{Cell} is strongly linearizable to v_{FLiT} , in the sense of [31]. This means that we assume that $v_{\text{FLiT}} \subseteq v'_{\text{Cell}}$ in addition to $v'_{\text{Cell}} \rightsquigarrow v_{\text{FLiT}}$. Note as well that the FLiT correctness theorem also assumes that a durably linearizable FLiT implementation is available, that is that an object $v'_{\text{FLiT}} : \dagger \text{FLiT}^{\sharp}$ durably linearizable to v_{FLiT} has been established. We do this for the implementation displayed in §1 in §1.

PROPOSITION E.16 (FLiT CORRECTNESS). For any object signature E , writing $v'_{\text{Mem}} := \otimes_{i \in I} v'_{\text{Cell}, i}$, if $v'_{\text{Mem}}; M$ is an object linearizable to v_E then, writing $v'_{\text{FLiTMem}} := \boxtimes_{i \in I} v'_{\text{FLiT}, i}$, it follows that $v'_{\text{FLiTMem}}; \text{dur}(M)$ is durably linearizable to v_E .

PROOF. Note that since v'_{Cell} is linearizable to v_{FLiT} , from locality for compositional linearizability it follows that $v'_{\text{Mem}} \rightsquigarrow v_{\text{Mem}}$, where we write $v_{\text{Mem}} := \otimes_{i \in I} v_{\text{FLiT}, i}$. Then, by observational refinement for compositional linearizability and the assumption, we have:

$$v_{\text{Mem}}; M = v'_{\text{Mem}}; M \subseteq v_E$$

Now, by locality for durable linearizability we have that $v'_{\text{FLiTMem}} \subseteq \text{dur}(v_{\text{Mem}})$. But then, the result follows from observational refinement for durable linearizability. \square

F IMPERATIVE PROGRAMS

So far we have developed the theory of crash-aware, strict and durable linearizability in a rather general setting. In practice, it proves useful to focus attention to strategies that specifically represent imperative code. Specifically, these are strategies that arise from parallel compositions of sequential imperative strategies. Object specifications $v_E : \dagger E^{\sharp}$ are now assumed to have effect signatures as types, but otherwise are just any strategy in the appropriate domain. Strategies $M : \dagger E^{\sharp} \multimap \dagger F^{\sharp}$ that are used to implement new objects $v_E; M$ will be specialized to these imperative strategies. For this, we use the theory of object-based semantics proposed by Reddy [36] and further developed in Oliveira Vale et al. [30, 31]. For brevity, we do this with respect to \mathbf{Crash} and \mathbf{Crash} , but the corresponding variation for durable strategies is easily obtained by enforcing durability throughout.

F.1 Parallel Strategies

Recall that in §A.1 we defined an operation $\text{Conc} -$ for constructing concurrent games from α -indexed collections of sequential games. This operation has a suitable counterpart for strategies, which makes it into a functor.

Definition F.1. Given a collection of strategies $\sigma[\Upsilon] = (\sigma[\alpha])_{\alpha \in \Upsilon}$, where for each α , $\sigma[\alpha] : A[\alpha] \multimap B[\alpha] \in \mathbf{Seq}$, define the strategy $\text{Conc } \sigma[\Upsilon] : \text{Conc } A \multimap \text{Conc } B$ as

$$\sigma[\Upsilon] := \parallel_{\alpha \in \Upsilon} \sigma[\alpha]$$

We say a strategy in the image of $\text{Conc} -$ is a *parallel strategy*.

Now, we define parallel strategies in **Crash**.

Definition F.2. We denote by $\text{Par} -$ the composition of semifunctors:

$$\text{Par} - : \mathbf{Seq}^\Upsilon \rightarrow \mathbf{Crash} := \text{vol}(\text{Conc} -)$$

We say a strategy in the image of Par is a crash-aware parallel strategy.

Since both $\text{Conc} -$ and $\text{Vol} -$ restrict to functors, we define the subsemicategory **Parallel** of **Crash** of parallel strategies, which, when restricted to saturated strategies, forms a subcategory **Parallel** of **Crash**.

We note that

PROPOSITION F.3.

$$\text{Par } \sigma[\Upsilon] \otimes \text{Par } \sigma'[\Upsilon] = \text{Par } (\sigma \otimes \sigma')[\Upsilon]$$

where $(\sigma \otimes \sigma')[\alpha] = \sigma[\alpha] \otimes \sigma'[\alpha]$.

Since, moreover, all the structural morphisms on **Crash** are parallel strategies, it follows that **Parallel** inherits the symmetric monoidal structure of **Crash**.

F.2 The Crash-Aware Replay Modality

Parallel strategies make for a nice domain for us to fully define the structure of the replay modality $\dagger -$, which we have been using for our examples.

Definition F.4. Given a sequential game $A = (M_A, \lambda_A, P_A)$ we define the sequential game $\dagger A$ by the following data:

$$M_{\dagger A} := \sum_{i \in \mathbb{N}} M_A \quad \lambda_{\dagger A} = \sum_{i \in \mathbb{N}} \lambda_A \quad P_{\dagger A} := \{s_1 \cdot \dots \cdot s_n \in \mathbb{P}_{\dagger A}^{\text{seq}} \mid \forall i. s_i \in \mathbf{i}:P_A\}$$

Given a sequential strategy $\sigma : A \multimap B$ we define $\dagger \sigma : \dagger A \multimap \dagger B$ by

$$\dagger \sigma := \{\mathbf{1}:s_1 \cdot \dots \cdot \mathbf{n}:s_n \in P_{\dagger A} \mid \forall i. s_i \in \sigma\}$$

Then, given a crash-aware game $A = (M_A, \lambda_A, P_A) \in \mathbf{Crash}$ we define the game $\dagger A$ by the following data.

$$M_{\dagger A} := \left(\sum_{i \in \mathbb{N}} M_A^\Upsilon \right) + \zeta \quad \lambda_{\dagger A} := \sum_{i \in \mathbb{N}} M_A \quad P_{\dagger A} := ((\parallel_{\alpha \in \Upsilon} P_{\dagger A}^\alpha) \cdot \zeta) \cdot (\parallel_{\alpha \in \Upsilon} P_{\dagger A}^\alpha)$$

Given a parallel strategy $\sigma = \text{Par } \sigma[\Upsilon]$ we define $\dagger \sigma := \text{Par } (\dagger \sigma[\alpha])_{\alpha \in \Upsilon}$.

A crucial result to appropriately define an object-based semantics model is that $\dagger -$ does define a modality, that is

PROPOSITION F.5.

$$\dagger - : \mathbf{Parallel} \rightarrow \mathbf{Parallel}$$

defines a semifunctor restricting to a comonad

$$\dagger - : \mathbf{Parallel} \rightarrow \mathbf{Parallel}$$

Object-based semantics then postulates that co-algebras of $\dagger -$ capture various flavors of the semantics of imperative code. It turns out that in fact, $\mathbf{Par} -$ transports co-algebras in

PROPOSITION F.6. *Given a collection of strategies $(M[\alpha] : A[\alpha] \multimap B[\alpha])_{\alpha \in \Upsilon}$ such that, for every $\alpha \in \Upsilon$, $A[\alpha]$ and $B[\alpha]$ are co-algebras of the sequential $\dagger -$, and $M[\alpha]$ is a co-algebra morphism, then $\mathbf{Par} (M[\Upsilon], R)$ is a co-algebra morphism with respect to the crash-aware $\dagger -$.*

This rather technical result means that we may use the same notion of imperative code used in Oliveira Vale et al. [30, 31] in our framework, and, therefore, we are able to inherit many of their techniques for modeling imperative programming to our setting.

F.3 Imperative Strategies

We are finally ready to define our model of imperative strategies.

Definition F.7. We define the category \mathbf{Imp} to be the subcategory of $\mathbf{Parallel}$ defined by the following data:

Objects: games of the form $E^\ddagger \in \mathbf{Crash}$, where E is a concurrent effect signature.

Morphisms: parallel strategies of the form $\mathbf{Par} (\widehat{M}[\Upsilon], \Delta_\ddagger) : \dagger E^\ddagger \multimap \dagger F^\ddagger$, where, for each α , $M[\alpha] : \dagger E[\alpha] \multimap F[\alpha] \in \mathbf{Seq}$, and $\widehat{M}[\alpha]$ is the co-Kleisli extension $\widehat{M}[\alpha] : \dagger E \multimap \dagger F$ of $M[\alpha]$.

This dense definition encapsulates a recent approach for the semantics of systems. The maps $M[\alpha] : \dagger E[\alpha] \multimap F[\alpha]$ are exactly the regular maps of Oliveira Vale et al. [30], which they show are effective at describing sequential imperative code. Our parallel strategies, in each epoch, play as parallel compositions of regular maps, which is the same notion of imperative concurrent code used in Oliveira Vale et al. [31]. Both of these trace back to the foundational work by Reddy [35, 36].

Two crucial results for the semantics of imperative programs are that:

PROPOSITION F.8. *For concurrent effect signatures E and F , the game $E^\ddagger \& F^\ddagger$ is generated by the concurrent effect signature $(E[\alpha] + F[\alpha])_{\alpha \in \Upsilon}$.*

PROPOSITION F.9 (IMPERATIVE SEELY ISOMORPHISM). *There is a natural isomorphism in \mathbf{Imp} :*

$$\dagger(E^\ddagger \& F^\ddagger) \cong \dagger E^\ddagger \otimes \dagger F^\ddagger$$

These two results together mean that \mathbf{Imp} inherits the symmetric monoidal structure of \mathbf{Vol} . Essentially, a tensor product $\dagger E^\ddagger \otimes \dagger F^\ddagger$ is essentially the same thing as the game $\dagger(E^\ddagger \& F^\ddagger)$, which, since $E^\ddagger \& F^\ddagger$ is generated by a concurrent effect signature, itself belongs to \mathbf{Imp} .

G A PROGRAM LOGIC FOR DURABLE OVERLAY OBJECTS

G.1 Programming Language

In this section, we define a general programming language. Compared to the main text, we define in detail the state transformer $\llbracket B \rrbracket_\alpha$ and the local operational semantics that we lift the transformer into.

G.1.1 Syntax. We start by defining a language Com for commands over some effect signature $E \in \text{Eff}$, where Eff is the set of effect signatures:

$$\text{Prim} := x \leftarrow e(a) \mid \text{assume}(\phi) \mid \text{ret } v \quad \text{Com} := \text{Prim} \mid \text{Com}; \text{Com} \mid \text{Com} + \text{Com} \mid \text{Com}^* \mid \text{skip}$$

Prim stands for primitive commands. The assignment command, $x \leftarrow e(a)$, executes the effect $e \in E$ with argument a and stores the response to variable x in a local environment $\Delta \in \text{Env}$. The assert command, $\text{assume}(\phi)$, takes a boolean function ϕ over the local environment and terminates the computation if it evaluates to False. We implement while loops and if-statements using $\text{assume}(-)$ in the usual way. The return command, $\text{ret } v$, stores the value v into a reserved variable res , and may only be invoked once in any procedure's execution. Com is the grammar of commands defined as usual in a Kleene algebra.

An implementation $M[\alpha]$ of type $E \rightarrow F \cup R_F$ implements the overlay's regular procedures F and recovery procedures R_F , using the underlay with the signature E . For simplicity, we require that there is only one recovery program in R_F , i.e. $R_F = \{r : \mathbf{1} \rightarrow \mathbf{1}\}$, and use r to denote the overlay's recovery method. The local implementation consists of a collection $M[\alpha] = (M[\alpha]^f)_{f \in F \cup R_F}$ of commands $M[\alpha]^f \in \text{Com}$ indexed by $f \in F \cup R_F$. A concurrent module $M[\Upsilon] \in \text{CMod}$ is given by a collection of local implementations $M[\Upsilon] = (M[\alpha])_{\alpha \in \Upsilon}$.

G.1.2 Semantics. Each primitive command B receives an interpretation as a state transformer $\llbracket B \rrbracket_\alpha : \text{UndState} \rightarrow \mathcal{P}(\text{UndState})$ over a set of states $\text{UndState} := \text{Env} \times P_{\dagger(E)}$ and returning a new set of states. A state $(\Delta, s) \in \text{UndState}$ contains a local environment $\Delta \in \text{Env}$ and a history represented as a play $s \in P_{\dagger(E)}$. The transformer $\llbracket B \rrbracket_\alpha$ depends on α only in that it tags each event it adds to the history with an agent identifier α . We define the transformer $\llbracket B \rrbracket_\alpha$ as follows.

- A special primitive is **id**, which is generated when reducing structural commands. Therefore, it has no effect on the state.

$$\llbracket \text{id} \rrbracket_\alpha(\Delta, s) = \{(\Delta, s)\}$$

- The return instruction has the interpretation below.

$$\llbracket \text{ret } v \rrbracket_\alpha(\Delta, s) = \begin{cases} \{(\Delta[\text{res} \mapsto v], s)\} & \Delta(\text{res}) = \perp \\ \emptyset & \Delta(\text{res}) \neq \perp \end{cases}$$

- For the assignment, we interpret it differently according to the underlay's trace.
 - If $\text{even}(\pi_\alpha(s))$, then

$$\llbracket x \leftarrow e(a) \rrbracket_\alpha(\Delta, s) = \begin{cases} \{(\Delta, s \cdot \alpha:e(a))\} & a \in \text{par}(e) \\ \{(\Delta, s \cdot \alpha:e(\Delta(a)))\} & a \in \text{Var} \wedge \Delta(a) \in \text{par}(e) \\ \emptyset & \text{otherwise} \end{cases}$$

- If $\pi_\alpha(s) = p \cdot e(a')$, where either $a' = a$ or $a' = \Delta(a)$, then

$$\llbracket x \leftarrow e(a) \rrbracket_\alpha(\Delta, s) = \{(\Delta[x \mapsto v], s \cdot \alpha:v \mid v \in \text{ar}(e))\}$$

- Otherwise, $\llbracket x \leftarrow e(a) \rrbracket_\alpha(\Delta, s) = \emptyset$.

- The assert instruction only makes progress when the boolean expression evaluates to true.

$$\llbracket \text{assume}(\phi) \rrbracket_\alpha(\Delta, s) = \begin{cases} \{(\Delta, s)\} & \phi(\Delta) = \text{True} \\ \emptyset & \text{otherwise} \end{cases}$$

We lift the interpretation function to a thread local operational semantics $\langle C, \Delta, s \rangle \xrightarrow{\alpha} \langle C', \Delta', s' \rangle$. It encodes how α steps on commands in a mostly standard way following the Kleene algebra structure of commands. We define the local operational semantics in Fig. 10.

$$\begin{array}{c}
\mapsto \subseteq \text{Com} \times \text{Prim} \times \{O, P\} \times \text{Com} \\
\\
\frac{}{B \mapsto_B^O B} \quad \frac{}{B \mapsto_B^P \text{skip}} \quad \frac{C_1 \mapsto_B^X C'_1}{C_1; C_2 \mapsto_B C'_1; C_2} \quad \frac{}{\text{skip}; C \mapsto_{\text{id}}^X C} \quad \frac{}{C^* \mapsto_{\text{id}}^X C; C^*} \\
\\
\frac{}{C^* \mapsto_{\text{id}}^X \text{skip}} \quad \frac{}{C_1 + C_2 \mapsto_{\text{id}}^X C_1} \quad \frac{}{C_1 + C_2 \mapsto_{\text{id}}^X C_2} \\
\\
\longrightarrow \subseteq (\text{Com} \times \text{UndState}) \times \Upsilon \times (\text{Com} \times \text{UndState}) \\
\\
\frac{(\Delta', s') \in \llbracket B \rrbracket_{\alpha}^X(\Delta, s) \quad C \mapsto_B^X C'}{\langle C, \Delta, s \rangle \longrightarrow_{\alpha} \langle C', \Delta', s' \rangle}
\end{array}$$

Fig. 10. Local Operational Semantics (\longrightarrow)

In Fig. 5, we lift this local operational semantics to a concurrent module operational semantics $\langle c, \Delta, s \rangle \longrightarrow_{R_E}^M \langle c', \Delta', s' \rangle$, which takes a continuation $c \in \text{Cont} := \Upsilon \rightarrow \{\text{idle}, \text{skip}, \text{dead}, \text{halt}\} + \text{Com}$ and a module state $(\Delta, s) \in \text{ModState} := (\Upsilon \rightarrow \text{Env}) \times P_{\dagger(E \cup R_E) \rightarrow \dagger(F \cup R_F)}$ containing local environments for all agents and the global trace of the system. It adds three highlighted rules to handle crashes, compared with the semantics in Oliveira Vale et al. [31]. These rules are:

- CRASH** Allows for crashes to happen at any time, resetting local environments for all agents, marking all the previously ran agents as dead and all remaining ones as halt.
- STARTREC** Starts the recovery phase by putting a recovery code C as the continuation, which is a sequential composition of one permutation of underlay recoveries followed by the overlay recovery $M[\alpha]^r$. Chajed et al. [8] uses a similar recovery scheme.
- ENDRED** When the recovery finishes, any thread that is not dead becomes idle, that the system can now run normally. It ensures the durable assumption since threads in previous epochs are no longer available.

We define the denotation of a module to be the set of traces it can generate under the module operational semantics from the initial configuration by the formula below, where c_0 is the initial continuation and Δ_0 is the initial environment where every agent has an empty local environment.

$$\llbracket M \rrbracket_{R_E} := \{s \mid \exists c \in \text{Cont}, \Delta \in (\Upsilon \rightarrow \text{Env}). \langle c_0, \Delta_0, \epsilon \rangle \longrightarrow_{R_E}^M \langle c, \Delta, s \rangle\} \subseteq P_{\dagger(E \cup R_E) \rightarrow \dagger(F \cup R_F)}$$

G.2 Object Interfaces

The interface of a crash-aware linearizable object E is a tuple

$$(v'_E : \dagger(E \cup R_E) \in \text{Dur}, v_E : \dagger E^\ddagger \in \text{Crash}) \quad \text{s.t.} \quad v'_E \upharpoonright_{E^\ddagger} \subseteq K_\ddagger v_E$$

where v'_E is the concrete specification that contains all possible traces the object can produce, which will include concurrent ones and will also contain crash events and recovery signatures, and v_E is the linearized specification. The interface is valid if and only if after recovery refining the concrete specification (the projection onto E^\ddagger), $v'_E \upharpoonright_{E^\ddagger}$ is crash-aware linearizable to v_E , i.e., $v'_E \upharpoonright_{E^\ddagger} \subseteq K_\ddagger v_E$.

Similarly, we define the interface of a durable linearizable object E as a tuple

$$(v'_E : \dagger(E \cup R_E) \in \text{Dur}, v_E : \dagger E \in \text{Conc}) \quad \text{s.t.} \quad v'_E \upharpoonright_{E^\ddagger} \subseteq \text{dur}(v_E)$$

A major difference from the crash-aware interface is that the durable interface's linearized specification v_E does not have crashes, because durable objects can be used as if there is no crash. The interface is valid if and only if $v'_E \upharpoonright_{E^\ddagger}$ is durable linearizable to v_E , i.e., $v'_E \upharpoonright_{E^\ddagger} \subseteq \text{dur}(v_E)$.

$$\begin{array}{c}
\longrightarrow \subseteq (\text{Com} \times \text{UndState}) \times \Upsilon \times (\text{Com} \times \text{UndState}) \\
\longrightarrow_{R_E} \subseteq (\text{Cont} \times \text{ModState}) \times \text{CMod} \times (\text{Cont} \times \text{ModState}) \\
\frac{f \in F \quad a \in \text{par}(f) \quad \Delta' = \Delta[\alpha \mapsto [\text{arg} \mapsto a]]}{\langle c[\alpha \mapsto \text{idle}], \Delta, s \rangle \longrightarrow_{R_E}^M \langle c[\alpha \mapsto M[\alpha]^f], \Delta', s \cdot \alpha : f \rangle} \text{INV} \\
\frac{\langle C, \Delta, s \upharpoonright_E \rangle \longrightarrow_\alpha \langle C', \Delta', s' \upharpoonright_E \rangle}{\langle c[\alpha \mapsto C], \Delta, s \rangle \longrightarrow_{R_E}^M \langle c[\alpha \mapsto C'], \Delta', s' \rangle} \text{STEP} \\
\frac{\pi_\alpha(s \upharpoonright_F) = p \cdot f \quad \Delta(\alpha)(\text{res}) = v \in \text{ar}(f) \quad \Delta' = \Delta[\alpha \mapsto \emptyset]}{\langle c[\alpha \mapsto \text{skip}], \Delta, s \rangle \longrightarrow_{R_E}^M \langle c[\alpha \mapsto \text{idle}], \Delta', s \cdot \alpha : v \rangle} \text{RET} \\
\frac{\forall \alpha \in s.c'[\alpha] = \text{dead} \\ \forall \alpha \in \Upsilon. \alpha \notin s \Rightarrow c'[\alpha] = \text{halt}}{\langle c, \Delta, s \rangle \longrightarrow_{R_E}^M \langle c', \Delta_0, s \cdot \downarrow \rangle} \text{CRASH} \\
\frac{s = s' \cdot \downarrow \quad \vec{r} = \text{perm}(R_E) \\ C = \text{sequence}(\vec{r}, M[\alpha]^r)}{\langle c[\alpha \mapsto \text{halt}], \Delta, s \rangle \longrightarrow_{R_E}^M \langle c[\alpha \mapsto C], \Delta, s \cdot \alpha : r \rangle} \text{STARTREC} \\
\frac{\pi_\alpha(s \upharpoonright_{F \cup R_F}) = s' \cdot r \quad \Delta(\alpha)(\text{res}) = v \in \text{ar}(r) \quad \Delta' = \Delta[\alpha \mapsto \emptyset] \\ \forall \alpha \in \Upsilon. c[\alpha] = \text{dead} \Rightarrow c'[\alpha] = \text{dead} \quad \forall \alpha \in \Upsilon. c[\alpha] \neq \text{dead} \Rightarrow c'[\alpha] = \text{idle}}{\langle c[\alpha \mapsto \text{skip}], \Delta, s \rangle \longrightarrow_{R_E}^M \langle c', \Delta', s \cdot \alpha : v \rangle} \text{ENDREC} \\
\text{where } \text{sequence}(\vec{r}, C) = \begin{cases} C & \vec{r} = \epsilon \\ (x_r \leftarrow r(a)); \text{sequence}(\vec{r}', C) & \vec{r} = r \cdot \vec{r}' \wedge a \in \text{par}(r) \wedge \text{reserved}(x_r) \end{cases}
\end{array}$$

Fig. 11. Module Operational Semantics (\longrightarrow_{R_E})

Usually, the client of the overlay object F will follow certain constraints when using the it. For example, when using a lock, one is supposed to invoke the lock acquire and the lock release in an alternating fashion. These constraints on clients often help us to prove a stronger and more useful specification v_F . We use a strategy $\mu_F : \dagger(F \cup R_F)$ to encode these client specifications. In the main text, we do not consider the client specification due to space limit, and we consider it in the program logic in this appendix. Most of the proof rules are the same except the PRIM rule.

The objective of our program logic is to establish the judgement

$$\mu_F \vdash M : (v'_E, v_E) \rightarrow (v'_F, v_F) \quad \text{or} \quad \mu_F \vdash M : (v'_E, v_E) \rightarrow \langle v'_F, v_F \rangle$$

which means under the assumption that the client will use the overlay F according to strategy μ_F , the implementation M implements F with either a crash-aware interface (v'_F, v_F) or a durable interface $\langle v'_F, v_F \rangle$, using the crash-aware underlay E with a valid interface (v'_E, v_E) . Concrete specification v'_F is defined by running M above v'_E of the underlay, i.e., $v'_F = (v'_E; \llbracket M \rrbracket_{R_E} \cap \mu_F) \upharpoonright_{(F \cup R_F)}$. The program logic's soundness guarantees the validity of the crash-aware/durable overlay interface. With the

validity, we may use the object F and its interface to implement and verify another layer of objects above it.

G.3 The Rely-Guarantee Crash Linearizability Hoare Logic (CLHL) for Durable Linearizability

We have been using a simplified rely-guarantee crash Hoare logic in the main text, while we develop a more expressive one in this appendix. Their main difference is that the CLHL in this appendix uses a binary relation between the pre-state and post-state of a program as the post-condition. This gives the logic more expressiveness and allows us to verify more complicated programs. We prove the CLHL with binary post-conditions to be sound, which implies the one with unary post-conditions in the main text to be sound, because it is strictly less expressive than the former one.

The program logic uses as proof configurations triples $(\Delta, s, \rho) \in \text{Config} := \text{ModState} \times \text{Poss}$, where Poss is a set of possibilities and is of type $\dagger F$. We define a configuration triple to be valid if and only if $s \upharpoonright_F$ is linearizable to ρ and ρ is linearizable to v_F . This is exactly the definition of the durable linearizability: after removing recoveries and crashes, the trace is linearizable to its specification. We maintain this as an invariant in proofs to ensure that the concrete trace s is always durably linearizable to v_F after the recovery refinement. Pre-conditions P and crash post-conditions Q_{\ddagger} are given by sets of configurations, while post-conditions Q , rely conditions \mathcal{R} , and guarantee conditions \mathcal{G} are specified as relations over the configurations. As usual, we define the stability requirements:

$$\text{stable}(\mathcal{R}, P) \iff \mathcal{R} \circ P \subseteq P \quad \text{stable}(\mathcal{R}, Q) \iff \mathcal{R} \circ Q \subseteq Q \wedge Q \circ \mathcal{R} \subseteq Q$$

G.3.1 Top Level Rules. The top level rule **OBJECT IMPL** proves M implements the overlay $\langle v'_F, v_F \rangle$ using the underlay (v'_E, v_E) .

$$\frac{\begin{array}{l} \forall \alpha, \alpha' \in \Upsilon. \alpha \neq \alpha' \Rightarrow \mathcal{G}[\alpha] \cup \text{invoke}_{\alpha}(-) \cup \text{return}_{\alpha}(-) \subseteq \mathcal{R}[\alpha'] \\ \forall \alpha \in \Upsilon. \mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \vDash_{\alpha}^F M[\alpha] \quad \forall \alpha \in \Upsilon. I \vDash_{\alpha}^R M[\alpha] \end{array}}{\mu F \vdash M : (v'_E, v_E) \rightarrow \langle v'_F, v_F \rangle} \text{ OBJECT IMPL}$$

It requires the prover to find an object invariant $I : \Upsilon \rightarrow \text{Config} \rightarrow \text{Prop}$ for the implementation and then prove the correctness of regular procedures and the recovery separately:

- *Verifying Regular Procedures.* To verify a concurrent object, the **OBJECT IMPL** rule requires finding the rely \mathcal{R} and guarantee \mathcal{G} of the object. The rely $\mathcal{R}[\alpha']$ of an agent needs to consider the effect of any other thread's executions, invocations, and returns, each represented by actions in $\mathcal{G}[\alpha]$, $\text{invoke}_{\alpha}(-)$, and $\text{return}_{\alpha}(-)$. And provers need to show $\mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \vDash_{\alpha}^F M[\alpha]$, which asserts that running regular procedures in F on the thread α , its environment will restrict its behavior in $\mathcal{R}[\alpha]$ while itself will only have behaviors in $\mathcal{G}[\alpha]$, and $I[\alpha]$ is satisfied when the thread α is idle.

Auxilliary Relations. We define some useful auxiliary relations here.

$$\begin{aligned} (\Delta, s, \rho) \text{invoke}_{\alpha}(f)(\Delta', s', \rho') &\iff \left((\Delta, s, \rho) \in \text{idle}_{\alpha} \wedge s' \upharpoonright_F \in \mu F \wedge \exists a. \Delta'(\alpha) = [\arg \mapsto a] \wedge \right. \\ &\quad \left. \forall \alpha' \neq \alpha. \Delta'(\alpha') = \Delta(\alpha) \wedge s' = s \cdot \alpha : f \wedge \rho' = \rho \cdot \alpha : f \right) \\ (\Delta, s, \rho) \text{return}_{\alpha}(f)(\Delta', s', \rho') &\iff \left((\Delta', s', \rho') = (\Delta, s, \rho) \wedge \right. \\ &\quad \left. \exists v \in \text{ar}(f). \Delta(\alpha)(\text{ret}) = v \wedge \text{last}(\pi_{\alpha}(\rho)) = \alpha : v \right) \\ (\Delta, s, \rho) \text{return}_{\alpha}(f(a))(\Delta', s', \rho') &\iff \left(\exists v \in \text{ar}(f). \Delta(\alpha)(\text{ret}) = v \wedge \Delta' = \emptyset \wedge \right. \\ &\quad \left. \rho' = \rho \wedge \text{last}(\pi_{\alpha}(\rho)) = \alpha : v \wedge s' = s \cdot \alpha : v \right) \end{aligned}$$

- The invoke relation requires the current thread α is idle in the pre-state, i.e., there is no pending invocation by α . The invoke relation then initialize the local environment according to method arguments and append the invocation event to both concrete trace s and possibility ρ . It also requires that after this invocation, s' satisfies the client specification. By composing this relation to pre-conditions, $\text{invoke}_\alpha(f) \circ P$, the result ignores traces where client specification is violated, since both programmers and provers have no obligations to ensure the correctness in that case.
- The returned relation requires that the post-state's local environment already has the reserved variable `ret` assigned some value and the possibility contains the response to the latest invocation. By composing it to some post-condition, $\text{returned}(f) \circ Q$, it requires provers to show that the latest invocation has returned and gets linearized.
- The return relation is a subsequent operation of the returned. It finishes the latest invocation by clearing its local environment and append the repose to the concrete trace s . The invoke and return are necessary because the program itself cannot manipulate the concrete trace by appending events of the overlay object. They are added by another client. And the invoke and return play the role of these clients by appending invocations to form a correct pre-condition and appending response to truly end a method execution.

Provers need to show $\mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \vDash_\alpha M_F[\alpha]$, which asserts that running the local implementation on thread α by invoking its methods, its environment will restrict their behaviors in $\mathcal{R}[\alpha]$ while itself will only have behaviors in $\mathcal{G}[\alpha]$, and $I[\alpha]$ is satisfied when the thread α is idle.

$$\forall \alpha, \alpha' \in \Upsilon. \alpha \neq \alpha' \Rightarrow \mathcal{G}[\alpha] \cup \text{invoke}_\alpha(-) \cup \text{return}_\alpha(-) \subseteq \mathcal{R}[\alpha'] \quad \forall \alpha \in \Upsilon. \mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \vDash_\alpha M_F[\alpha]$$

The LOCAL IMPL rule proves this judgement by splitting $I[\alpha]$ into conjunctions of $P[\alpha]^f$, each specifying the pre-condition of a method invocation, and then proving a series of objectives.

$$\frac{I[\alpha] = \bigcap_{f \in F} P[\alpha]^f \quad \forall f \in F. (\Delta_0, \epsilon, \epsilon) \in P[\alpha]^f \quad \forall f \in F. \text{stable}(\mathcal{R}[\alpha], P[\alpha]^f) \quad \forall f \in F. \mathcal{R}[\alpha], \mathcal{G}[\alpha] \vDash_\alpha^f \{P[\alpha]^f\} M[\alpha]^f \{Q[\alpha]^f\} \{\top\} \quad \forall f \in F. \text{return}_\alpha(f) \circ Q[\alpha]^f \subseteq I[\alpha]}{\mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \vDash_\alpha M_F[\alpha]} \quad \text{LOCAL IMPL}$$

- First of all, each pre-condition $P[\alpha]^f$ needs to include the initial configuration. Each pre-condition must be stable under interferences (the rely $\mathcal{R}[\alpha]$) of the environment, and therefore the invariant $I[\alpha]$ is also stable w.r.t. environment interfaces.
- Then, provers need to show that each method f satisfies

$$\mathcal{R}[\alpha], \mathcal{G}[\alpha] \vDash_\alpha^f \{P[\alpha]^f\} M[\alpha]^f \{Q[\alpha]^f\} \{\top\}$$

which is a shorthand for the CLHL hexad below by hiding auxiliary relations applied to pre-/post-conditions.

$$\mathcal{R}[\alpha], \mathcal{G}[\alpha] \vDash_\alpha \{\text{invoke}_\alpha(f) \circ P[\alpha]^f\} M[\alpha]^f \{\text{returned}_\alpha(f) \circ Q[\alpha]^f\} \{\top\}$$

A hexad of the form $\mathcal{R}, \mathcal{G} \vDash_\alpha \{P\} C \{Q\} \{Q_i\}$ means that given states satisfying P , running the program C in an environment with interference in \mathcal{R} and on thread α will produce actions in \mathcal{G} , and if it terminates normally, the state will satisfy Q , and if it crashes, the state will satisfy Q_i . A hexad is provable with proof rules introduced later. It is worth mentioning that *there is no need to explicitly specify and prove a crash post-condition for each regular method*, and we can simply put \top as the crash post-condition. This is true because:

- (1) The guarantee $\mathcal{G}[\alpha]$ of the current thread is included in any other thread's rely $\mathcal{R}[\alpha']$, and therefore any step during the execution of any method in thread α is captured in $\mathcal{R}[\alpha']$.
- (2) For any other thread α' , its invariant $I[\alpha']$ is stable w.r.t. $\mathcal{R}[\alpha']$, which means any state after any execution step of any method in thread α (captured in $\mathcal{R}[\alpha']$) is in $I[\alpha']$.

- (3) Therefore, the state of thread α will satisfy any other thread's invariant $I[\alpha']$ at any time (including the point of crash), and the crash post-condition in α can be derived from $I[\alpha']$.
- Lastly, after finished the execution of a method and returned from it, the program state need to satisfy the invariant so that the current thread can still access the object and invoke its procedures.

$$\forall f \in F. \text{return}_\alpha(f) \circ Q[\alpha]^f \subseteq I[\alpha]$$

- *Verifying the Recovery.* Then, to ensure the durability of the object, provers need to show $I \vDash_\alpha^R M[\alpha]$, which means whenever crash happens, the execution of the recovery on any thread α can restore the program state to satisfy the object invariant I . It can be verified via the RECOVER IMPL rule.

$$\frac{\text{ID}, \top \vDash_\alpha^r \{P_r[\alpha]\}M[\alpha]^r\{Q_r[\alpha]\}\{Q_\zeta[\alpha]\} \quad Q_\zeta[\alpha] \subseteq P_r[\alpha] \quad \cup_{\alpha' \in \Upsilon} I[\alpha'] \Rightarrow_\zeta Q_\zeta[\alpha] \quad \text{return}_\alpha(r) \circ Q_r[\alpha] \subseteq \cap_{\alpha' \in \Upsilon} I[\alpha']}{I \vDash_\alpha^R M[\alpha]} \text{RECOVER IMPL}$$

First of all, provers need to find the pre-condition P_r , the post-condition Q_r , and the crash post-condition Q_ζ of the recovery program, and prove the following hexad³

$$\text{ID}, \top \vDash_\alpha^r \{P_r[\alpha]\}M[\alpha]^r\{Q_r[\alpha]\}\{Q_\zeta[\alpha]\}$$

which means running the recovery program $M[\alpha]^r$ on arbitrary thread α from states in $P_r[\alpha]$ will either recover the system into states in $Q_r[\alpha]$ or crash into states in $Q_\zeta[\alpha]$. Since the recovery program always starts execution after a crash, the crash post-condition Q_ζ needs to imply the recovery's pre-condition P_r .

As mentioned before, $I[\alpha']$ will serve as the crash post-condition of other threads. Therefore, we require that all $I[\alpha']$ crash into the crash post-condition Q_ζ of the recovery program to ensure that all possible crashes are considered. The crash-into relation (\Rightarrow_ζ) transforms the assertion I into Q_ζ , which means that any state satisfying P will satisfy Q_ζ immediately after a crash.

$$I \Rightarrow_\zeta Q_\zeta \iff \forall (\Delta, s, \rho) \in I. (\Delta_0, s \cdot \zeta, \rho) \in Q_\zeta$$

Lastly, after the execution of the recovery, the system is restored and ready to run and therefore, the program state after the recovery's return needs to imply the invariant $I[\alpha']$ of any thread α' .

G.3.2 CLHL Proof Rules for the Hexad. According to these top level rules, proofs of both the implementation and the recovery boil down to proofs of hexads like $\mathcal{R}, \mathcal{G} \vDash_\alpha \{P\}C\{Q\}\{Q_\zeta\}$. Figure 12 shows CLHL's proof rules for the hexad.

Among CLHL proof rules for the hexad, the core proof rule for proving the durable linearizability is the PRIM rule. Firstly, we need to show that the pre-/post-condition and the crash post-condition can crash into (\Rightarrow_ζ) the crash post-condition, because the crash can happen at any time, even after another crash.

$$P \Rightarrow_\zeta Q_\zeta \quad Q \circ P \Rightarrow_\zeta Q_\zeta \quad Q_\zeta \Rightarrow_\zeta Q_\zeta$$

Then, as any usual rely-guarantee logic, the pre-/post-condition needs to be stable w.r.t. the rely.

$$\text{stable}(\mathcal{R}, P) \quad \text{stable}(\mathcal{R}, Q)$$

³We reuse the concurrent CLHL for the sequential recovery program, so the rely and guarantee for it are ID and \top .

$$\begin{array}{c}
\frac{P \Rightarrow_{\downarrow} Q_{\downarrow} \quad Q \circ P \Rightarrow_{\downarrow} Q_{\downarrow} \quad Q_{\downarrow} \Rightarrow_{\downarrow} Q_{\downarrow}}{\text{stable}(\mathcal{R}, P) \quad \text{stable}(\mathcal{R}, Q) \quad \mathcal{G} \vdash_{\alpha} \{P\}B\{Q\}} \text{PRIM} \\
\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}B\{Q\}\{Q_{\downarrow}\} \\
\\
\frac{\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}C_1\{Q_1\}\{Q_{\downarrow}\} \quad \mathcal{R}, \mathcal{G} \vDash_{\alpha} \{Q_1 \circ P\}C_2\{Q_2\}\{Q_{\downarrow}\}}{\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}C_1; C_2\{Q_2 \circ Q_1\}\{Q_{\downarrow}\}} \text{SEQ} \\
\\
\frac{\text{stable}(\mathcal{R}, P) \quad P \Rightarrow_{\downarrow} Q_{\downarrow}}{\mathcal{R}, \text{ID} \vDash_{\alpha} \{P\}\text{skip}\{\text{ID}\}\{Q_{\downarrow}\}} \text{SKIP} \quad \frac{\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}C\{Q\}\{Q_{\downarrow}\} \quad Q \circ P \subseteq P}{\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}C^*\{Q\}\{Q_{\downarrow}\}} \text{ITER} \\
\\
\frac{\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}C_1\{Q\}\{Q_{\downarrow}\} \quad \mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}C_2\{Q\}\{Q_{\downarrow}\}}{\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}C_1 + C_2\{Q\}\{Q_{\downarrow}\}} \text{CHOICE} \\
\\
\frac{\text{stable}(\mathcal{R}', P') \quad \text{stable}(\mathcal{R}', Q') \quad Q_{\downarrow} \subseteq Q'_{\downarrow} \quad Q'_{\downarrow} \Rightarrow_{\downarrow} Q'_{\downarrow} \quad Q' \circ P' \Rightarrow_{\downarrow} Q'_{\downarrow}}{P' \subseteq P \quad Q \subseteq Q' \quad \mathcal{R}' \subseteq \mathcal{R} \quad \mathcal{G} \subseteq \mathcal{G}' \quad \mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}C\{Q\}\{Q_{\downarrow}\}} \text{CONSEQ} \\
\mathcal{R}', \mathcal{G}' \vDash_{\alpha} \{P'\}C\{Q'\}\{Q'_{\downarrow}\}
\end{array}$$

Fig. 12. Proof Rules for the CLHL Hexad

Lastly, we need to prove the commit rule $\mathcal{G} \vdash_{\alpha} \{P\}B\{Q\}$ for the primitive command B , which is directly defined by and provable by the semantics:

$$\begin{array}{c}
\forall(\Delta, s, \rho). s \upharpoonright_{F \cup R_F} \in \mu_F \wedge (\Delta, s, \rho) \in P \wedge \\
\mathcal{G} \vdash_{\alpha} \{P\}B\{Q\} \iff \left(\begin{array}{c} \forall(\Delta', s') \in \llbracket B \rrbracket_{\alpha}(\Delta, s) \cap v_E \Rightarrow s' \upharpoonright_{F \cup R_F} \in \mu_F \wedge \\ \exists \rho'. (\Delta, s, \rho)Q(\Delta', s', \rho') \wedge (\Delta, s, \rho)\mathcal{G}(\Delta', s', \rho') \wedge \rho \dashrightarrow \rho' \end{array} \right) \\
\text{where } \rho \dashrightarrow \rho' \iff \exists t_P \in (M_F^P)^* . \rho \cdot t_P \rightsquigarrow_{\dagger F} \rho'
\end{array}$$

The commit rule states that command B will update configurations in P by appending the corresponding event to the concrete trace, which may be the commitment point of some pending operations. To maintain the invariant that s is durably linearizable to ρ , the commit rule allows a ghost update, $\rho \dashrightarrow \rho'$, where provers can append several response events to ρ and rewrite it according to $\rightsquigarrow_{\dagger F}$ to obtain ρ' , a new possibility that s linearizes into. After the update made by the command B and the angelic update by the prover, the new configuration needs to satisfy the post-condition Q , and both updates need to be recorded in the guarantee.

To summarize, there are following steps to prove the durable linearizability of a concurrent object using CLHL:

- (1) Firstly, find the rely condition \mathcal{R} , the guarantee condition \mathcal{G} , and the invariant I , and use OBJECT IMPL rule to generate separate proof goals for verifying regular procedures and the recovery program.
- (2) The second step is to find pre-conditions and normal post-conditions of regular procedures and apply LOCAL IMPL to generate Hoare hexads for verifying regular procedures.
- (3) The third step is to find the pre-condition, normal post-condition, and the crash post-condition (crash invariant) of the recovery program and prove the Hoare hexad for it and show that the object invariant will crash into the crash invariant.

- (4) Lastly, prove hexads of each procedure and other side conditions generated by top level rules using CLHL's proof rules for hexads.

Remark. Notice that for the regular procedure verification, we do not use the crash post-condition (and instead always set it to \top), and for the recovery verification, we do not use the rely and guarantee condition because it is a sequential program. The another design choice is to use two different logics for them, each without the unnecessary part. But we find that their logics will have almost the same set of proof rules, and by unifying them into one rely-guarantee crash Hoare logic not only benefits our presentation but also simplifies the soundness proof.

Another benefit is that we can easily extend the CLHL for the recovery verification to concurrent recovery programs, where multiple recovery programs are running concurrently on multiple thread. The rely-guarantee condition in the Hoare hexad makes the program logic ready for the concurrent recovery verification. Moreover, since the regular procedure verification and the recovery verification are decoupled (with only an object invariant linking them together), we do not need to change the LOCAL IMPL rule when we extend the RECOVER IMPL rule to concurrent recovery programs.

G.4 Soundness

The program logic is justified by the following soundness theorem.

PROPOSITION G.1 (SOUNDNESS). *If $\mu_F \vdash M : (v'_E, v_E) \rightarrow \langle v'_F, v_F \rangle$ is provable, and (v'_E, v_E) is a valid underlay interface, and $v'_F = v'_E; \llbracket M \rrbracket_{R_E} \cap \mu_F$, then $\langle v'_F, v_F \rangle$ is a valid overlay interface with*

$$v'_F \upharpoonright_{F^\sharp} \subseteq \text{dur}(v_F).$$

To prove the soundness (proposition G.1), we extend methods in [23, 31] and establish the overlay interface's validity in four steps depicted in the formula below. (1) We first use the recovery refinement in §5.1 to remove the underlay's refinement. (2) Then, by the validity of the underlay interface, and observational refinement, we can use their specification in the execution of the overlay instead of using their concrete traces. (3) We use the linking lemma G.2 to integrate underlay's specification in overlay's denotation, which makes the next step easier. (4) The key step is the auxiliary soundness (lemma G.6), which establish the linearization from overlay's concrete traces to its specification.

$$\begin{aligned}
 v'_F \upharpoonright_{F^\sharp} &= (v'_E; \llbracket M \rrbracket_{R_E} \cap \mu_F) \upharpoonright_{F^\sharp} && \text{By Definition of } v'_F \\
 &= ((v'_E; \llbracket M \rrbracket_{R_E}) \upharpoonright_{E^\sharp} \multimap_{F \cup R_F} \cap \mu_F) \upharpoonright_{F^\sharp} \\
 &\subseteq (v'_E \upharpoonright_{E^\sharp}; \llbracket M \rrbracket_{\emptyset} \cap \mu_F) \upharpoonright_{F^\sharp} && (1) \text{ Recovery Refinement} \\
 &\subseteq (v_E; \llbracket M \rrbracket_{\emptyset} \cap \mu_F) \upharpoonright_{F^\sharp} && (2) \text{ Validity of } (v'_E, v_E) + \text{Obs. Ref.} \\
 &= (\llbracket \text{Link } v_E; M \rrbracket \cap \mu_F) \upharpoonright_{F^\sharp} && (3) \text{ Linking (lemma G.2)} \\
 &\subseteq \text{dur}(v_F) && (4) \text{ Auxiliary Soundness (lemma G.6)}
 \end{aligned}$$

To make the auxiliary soundness proof easier, we first embed the underlay's specification into an auxiliary module semantics by $- \twoheadrightarrow_{v_E}^M -$.

$$\twoheadrightarrow \subseteq (\text{Cont} \times \text{ModState}) \times (\text{CMod} \times \text{Crash}) \times (\text{Cont} \times \text{ModState})$$

$$\frac{f \in F \quad a \in \text{par}(f) \quad \Delta' = \Delta[\alpha \mapsto [\text{arg} \mapsto a]]}{\langle c[\alpha \mapsto \text{idle}], \Delta, s \rangle \twoheadrightarrow_{v_E}^M \langle c[\alpha \mapsto M[\alpha]^f], \Delta', s \cdot \alpha : f \rangle}$$

$$\frac{\langle C, \Delta, s \upharpoonright_E \rangle \longrightarrow_{\alpha} \langle C', \Delta', s' \upharpoonright_E \rangle \quad s' \upharpoonright_E \in v_E}{\langle c[\alpha \mapsto C], \Delta, s \rangle \twoheadrightarrow_{v_E}^M \langle c[\alpha \mapsto C'], \Delta', s' \rangle}$$

$$\frac{\pi_{\alpha}(s \upharpoonright_F) = p \cdot f \quad \Delta(\alpha)(\text{res}) = v \in \text{ar}(f) \quad \Delta' = \Delta[\alpha \mapsto \emptyset]}{\langle c[\alpha \mapsto \text{skip}], \Delta, s \rangle \twoheadrightarrow_{v_E}^M \langle c[\alpha \mapsto \text{idle}], \Delta', s \cdot \alpha : v \rangle}$$

$$\frac{\forall \alpha \in s.c'[\alpha] = \text{dead} \quad \forall \alpha \in \Upsilon.\alpha \notin s \Rightarrow c'[\alpha] = \text{halt} \quad s = s' \cdot \downarrow}{\langle c, \Delta, s \rangle \twoheadrightarrow_{v_E}^M \langle c', \Delta_0, s \cdot \downarrow \rangle} \quad \frac{}{\langle c[\alpha \mapsto \text{halt}], \Delta, s \rangle \twoheadrightarrow_{v_E}^M \langle c[\alpha \mapsto M[\alpha]^r], \Delta, s \cdot \alpha : r \rangle}$$

$$\frac{\pi_{\alpha}(s \upharpoonright_{F \cup R_F}) = s' \cdot r \quad \Delta(\alpha)(\text{res}) = v \in \text{ar}(r) \quad \Delta' = \Delta[\alpha \mapsto \emptyset] \quad \forall \alpha \in \Upsilon.c[\alpha] = \text{dead} \Rightarrow c'[\alpha] = \text{dead} \quad \forall \alpha \in \Upsilon.c[\alpha] \neq \text{dead} \Rightarrow c'[\alpha] = \text{idle}}{\langle c[\alpha \mapsto \text{skip}], \Delta, s \rangle \twoheadrightarrow_{v_E}^M \langle c', \Delta', s \cdot \alpha : v \rangle}$$

The only difference between $- \twoheadrightarrow_{v_E}^M -$ and $- \twoheadrightarrow_{\emptyset}^M -$ is the rule that lifts thread local reductions, where we ask steps by the underlay satisfies its specification v_E . We define the linking denotation $\llbracket \text{Link}_{v_E}; M \rrbracket : \dagger E^{\downarrow} \multimap \dagger(F \cup R_F)$ by the formula

$$\llbracket \text{Link}_{v_E}; M \rrbracket := \{s \mid \exists c \in \text{Cont}, \Delta \in (\Upsilon \rightarrow \text{Env}). \langle c_0, \Delta_0, \epsilon \rangle \twoheadrightarrow_{v_E}^M \langle c, \Delta, s \rangle\} \subseteq P_{\dagger E^{\downarrow} \multimap \dagger(F \cup R_F)}.$$

Lemma G.2 allows the transformation between the module denotation and the auxiliary linking denotation. Its proof is similar to the one in Oliveira Vale et al. [31].

LEMMA G.2 (LINKING). *For any $M \in \text{CMod}$ and given $v_E : \dagger E^{\downarrow}$, we have*

$$v_E; \llbracket M \rrbracket_{\emptyset} = \llbracket \text{Link}_{v_E}; M \rrbracket$$

LEMMA G.3. *For any c, Δ, s, M, v_E , if $\langle c_0, \Delta_0, \epsilon \rangle \twoheadrightarrow_{v_E}^M \langle c, \Delta, s \rangle$, then*

$$\text{last}(s) = \downarrow \iff \forall \alpha.c(\alpha) \in \{\text{dead}, \text{halt}\}.$$

PROOF. By discussing the last reduction step in $\langle c_0, \Delta_0, \epsilon \rangle \twoheadrightarrow_{v_E}^M \langle c, \Delta, s \rangle$. \square

Definition G.4 (Safety Judgement). We define the judgement $\text{safe}_{\alpha}(\mathcal{R}, \mathcal{G}, P_0, P, s, Q, Q_{\downarrow})$ inductively as follows:

$$\frac{\text{rely}(\mathcal{R}, P) \subseteq Q \circ P_0}{\text{safe}_{\alpha}(\mathcal{R}, \mathcal{G}, P_0, P, \text{skip}, Q, Q_{\downarrow})} \text{ DONE}$$

$$\frac{\forall C'. C \twoheadrightarrow_B^X C' \Rightarrow \exists Q'. (\mathcal{G} \vdash_{\alpha} \{\text{rely}(\mathcal{R}, P)\} B\{Q'\} \wedge \text{stable}(\mathcal{R}, Q' \circ \text{rely}(\mathcal{R}, P)) \wedge \text{safe}_{\alpha}(\mathcal{R}, \mathcal{G}, P_0, Q' \circ \text{rely}(\mathcal{R}, P), C', Q, Q_{\downarrow}) \wedge Q' \circ \text{rely}(\mathcal{R}, P) \Rightarrow_{\downarrow} Q_{\downarrow})}{\text{safe}_{\alpha}(\mathcal{R}, \mathcal{G}, P_0, P, C, Q, Q_{\downarrow})} \text{ STEP}$$

where $\text{rely}(\mathcal{R}, P) \triangleq P \cup \mathcal{R} \circ P$.

LEMMA G.5. For any $\mathcal{R}, \mathcal{G}, P, s, Q, Q_{\ddagger}$, if the quadruple $\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}s\{Q\}\{Q_{\ddagger}\}$ is provable, then the followings are true.

$$\text{stable}(\mathcal{R}, P) \quad P \Rightarrow_{\ddagger} Q_{\ddagger} \quad Q_{\ddagger} \Rightarrow_{\ddagger} Q_{\ddagger} \quad \text{safe}_{\alpha}(\mathcal{R}, \mathcal{G}, P, P, s, Q, Q_{\ddagger})$$

PROOF. By induction over the derivation tree of $\mathcal{R}, \mathcal{G} \vDash_{\alpha} \{P\}s\{Q\}\{Q_{\ddagger}\}$. \square

LEMMA G.6 (AUXILIARY SOUNDNESS). If the judgement $\mu_F \vdash M : (v'_E, v_E) \rightarrow \langle v'_F, v_F \rangle$ is provable and (v'_E, v_E) is a valid crash-aware underlay interface, then

$$(\llbracket \text{Link } v'_E; M \rrbracket \cap \mu_F) \upharpoonright_{F\ddagger} \subseteq \text{dur}(v_F).$$

PROOF. Since $\mu_F \vdash M : (v'_E, v_E) \rightarrow \langle v'_F, v_F \rangle$ is provable, by the OBJECT IMPL rule, there exists $\mathcal{R}, \mathcal{G}, I, P_r, Q_r, Q_{\ddagger}$ with the following conclusions.

$$\forall \alpha, \alpha' \in A. \alpha \neq \alpha' \Rightarrow \mathcal{G}[\alpha] \cup \text{invoke}_{\alpha}(-) \cup \text{return}_{\alpha}(-) \subseteq \mathcal{R}[\alpha] \quad (\text{H-RG})$$

$$\forall \alpha \in \Upsilon. \mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \vDash_{\alpha} M_F[\alpha] \quad (\text{H-TLQ})$$

$$\forall \alpha, \alpha' \in A. Q_{\ddagger}[\alpha] \subseteq P_r[\alpha'] \quad (\text{H-RPre})$$

$$\forall \alpha. I[\alpha] \Rightarrow_{\ddagger} Q_{\ddagger}[\alpha] \quad (\text{H-PC})$$

$$\forall \alpha, \alpha' \in A. \text{return}_{\alpha}(r) \circ \text{returned}_{\alpha}(r) \circ Q_r[\alpha] \circ \text{invoke}_{\alpha}(r) \circ P_r[\alpha] \subseteq I[\alpha'] \quad (\text{H-RPost})$$

$$\forall \alpha \in A. \text{ID}, \top \vDash_{\alpha} \{\text{invoke}_{\alpha}(r) \circ P_r[\alpha]\} M[\alpha]^r \{\text{returned}_{\alpha}(r) \circ Q_r[\alpha]\} \{Q_{\ddagger}[\alpha]\} \quad (\text{H-RQ})$$

By the LOCAL IMPL rule and (H-TLQ), there exists $P[\alpha]^f, Q[\alpha]^f$ with the following conclusion about thread local executions of each regular function.

$$\bigcap_{f \in F} P[\alpha]^f = I[\alpha] \quad (\text{H-Asrt})$$

$$\forall \alpha \in A, f \in F. (\Delta_0, \epsilon, \epsilon) \in P[\alpha]^f \quad (\text{H-FInit})$$

$$\forall \alpha \in A, f \in F. \mathcal{R}[\alpha], \mathcal{G}[\alpha] \vDash_{\alpha} \{\text{invoke}_{\alpha}(f) \circ P[\alpha]^f\} M[\alpha]^f \{\text{returned}_{\alpha}(f) \circ Q[\alpha]^f\} \{\top\} \quad (\text{H-FQ})$$

$$\forall \alpha \in A, f, f' \in F. \text{return}_{\alpha}(f) \circ \text{returned}_{\alpha}(f) \circ Q[\alpha]^f \circ \text{invoke}_{\alpha}(f) \circ P[\alpha]^f \subseteq P[\alpha]^{f'} \quad (\text{H-FPP})$$

$$\forall f \in F. \text{stable}(\mathcal{R}[\alpha], P[\alpha]^f) \quad (\text{H-Stb})$$

To prove $(\llbracket \text{Link } v_E; M \rrbracket \cap \mu_F) \upharpoonright_{F\ddagger} \subseteq \text{dur}(v_F)$, we only need to prove

$$\forall s \in \llbracket \text{Link } v_E; M \rrbracket \cap \mu_F. s \upharpoonright_{F\ddagger} \in \text{dur}(v_F)$$

which is equivalent to the following by definitions of auxiliary denotation and durable linearizability.

$$\forall c, \Delta, s. \langle c_0, \Delta_0, \epsilon \rangle \xrightarrow{M}_{v_E} \langle c, \Delta, s \rangle \wedge s \upharpoonright_{F\ddagger \& R_F} \in \mu_F \Rightarrow \exists \rho_F \in v_F. s \upharpoonright_F \twoheadrightarrow \rho_F$$

We generalize this proposition to the following one.

PROPOSITION G.7. For any c, Δ, s , if $\langle c_0, \Delta_0, \epsilon \rangle \xrightarrow{M}_{v_E} \langle c, \Delta, s \rangle$, then when $s \upharpoonright_{F^\ddagger \& R_F} \in \mu_F$ there exists a current linearization $\rho_F \in \nu_F$ and the followings hold.

$$s \upharpoonright_F \dashv\dashv \rho_F \quad (\text{G-Lin})$$

$$\wedge \forall \alpha. c[\alpha] \neq \text{dead} \Rightarrow$$

$$\exists P_\alpha. (\Delta, s, \rho_F) \in P_\alpha \wedge (\text{halt} \notin c \Rightarrow \text{stable}(\mathcal{R}[\alpha], P_\alpha)) \quad (\text{G-Pa})$$

$$\wedge (c[\alpha] = \text{idle} \Rightarrow P_\alpha \subseteq I[\alpha]) \quad (\text{G-Idle})$$

$$\wedge ((\forall \alpha'. c[\alpha'] \in \{\text{dead}, \text{halt}\}) \Rightarrow \exists \alpha'. P_\alpha \subseteq Q_\ddagger[\alpha']) \quad (\text{G-Crs})$$

$$\wedge (c[\alpha] = \text{halt} \wedge (\exists \alpha', C \in \text{Com.c}[\alpha'] = C) \Rightarrow \exists \alpha'. P_\alpha \Rightarrow_\ddagger Q_\ddagger[\alpha']) \quad (\text{G-Rec})$$

$$\wedge \left(\forall C \in \text{Com.c}[\alpha] = C \Rightarrow \exists f \in F \cup R_F. \left(\begin{array}{l} \text{safe}_\alpha \left(\begin{array}{l} \mathcal{R}[\alpha]^f, \mathcal{G}[\alpha]^f, \text{invoke}_\alpha(f) \circ P[\alpha]^f \\ P_\alpha, C, \text{returned}_\alpha(f) \circ Q[\alpha]^f, Q_\ddagger[\alpha]^f \end{array} \right) \\ \wedge (f \in R_F \Rightarrow \forall \alpha' \neq \alpha. c(\alpha') \in \{\text{dead}, \text{halt}\}) \\ \wedge P_\alpha \Rightarrow_\ddagger Q_\ddagger[\alpha]^f \wedge f = \text{last}(\pi_\alpha(s \upharpoonright_{F^\ddagger \& R_F})) \end{array} \right) \right) \quad (\text{G-Exe})$$

For conciseness, We define $(P[\alpha]^r, Q[\alpha]^r, Q_\ddagger[\alpha]^r)$ as $(P_r[\alpha], Q_r[\alpha], Q_\ddagger[\alpha])$, and $Q_\ddagger[\alpha]^f = \top$ for $f \in F$, and $(\mathcal{R}[\alpha]^f, \mathcal{G}[\alpha]^f)$ to be $(\mathcal{R}[\alpha], \mathcal{G}[\alpha])$ for $f \in F$, and $(\mathcal{R}[\alpha]^r, \mathcal{G}[\alpha]^r)$ to be (ID, \top) for the recovery r . For the right conjunct of (G-Pa), we would usually have $\text{halt} \notin c$ and use/prove it simply as $\text{stable}(\mathcal{R}[\alpha], P_\alpha)$ at most of the time. We only consider the LHS when necessary.

To prove the auxiliary soundness, it suffice to prove proposition G.7 under hypotheses introduced so far. We prove it by induction on the length of the reduction $\langle c_0, \Delta_0, \epsilon \rangle \xrightarrow{M}_{v_E} \langle c, \Delta, s \rangle$.

- **Base Case.** In the base case, we have $(c_0, \Delta_0, \epsilon) = (c, \Delta, s)$. And we take the current linearization $\rho_F = \epsilon$, which satisfies (G-Lin). For each α , we take $P_\alpha = I[\alpha] = \bigcap_{f \in F} P[\alpha]^f$ ((H-Asrt) and (G-Pa)) is apparent by (H-FInit) and (G-Idle). Other conditions (G-Exe), (G-Rec), and (G-Crs) because the LHS of the implication is false.

- **Inductive Step.** In the inductive step, we have

$$\langle c_0, \Delta_0, \epsilon \rangle \xrightarrow{M^*}_{v_E} \langle c, \Delta, s \rangle \text{ and } \langle c, \Delta, s \rangle \xrightarrow{M}_{v_E} \langle c', \Delta', s' \rangle$$

and the induction hypothesis that if $s \upharpoonright_{F^\ddagger \& R_F} \in \mu_F$, then there exists $\rho_F \in \nu_F$ with the followings.

$$s \upharpoonright_F \dashv\dashv \rho_F \quad (\text{IH-Lin})$$

$$\wedge \forall \alpha. c[\alpha] \neq \text{dead} \Rightarrow$$

$$\exists P_\alpha. (\Delta, s, \rho_F) \in P_\alpha \wedge (\text{halt} \notin c \Rightarrow \text{stable}(\mathcal{R}[\alpha], P_\alpha)) \quad (\text{IH-Pa})$$

$$\wedge (c[\alpha] = \text{idle} \Rightarrow P_\alpha \subseteq I[\alpha]) \quad (\text{IH-Idle})$$

$$\wedge ((\forall \alpha'. c[\alpha'] \in \{\text{dead}, \text{halt}\}) \Rightarrow \exists \alpha'. P_\alpha \subseteq Q_\ddagger[\alpha']) \quad (\text{IH-Crs})$$

$$\wedge (c[\alpha] = \text{halt} \wedge (\exists \alpha', C \in \text{Com.c}[\alpha'] = C) \Rightarrow \exists \alpha'. P_\alpha \Rightarrow_\ddagger Q_\ddagger[\alpha']) \quad (\text{IH-Rec})$$

$$\wedge \left(\forall C \in \text{Com.c}[\alpha] = C \Rightarrow \exists f \in F \cup R_F. \left(\begin{array}{l} \text{safe}_\alpha \left(\begin{array}{l} \mathcal{R}[\alpha]^f, \mathcal{G}[\alpha]^f, \text{invoke}_\alpha(f) \circ P[\alpha]^f \\ P_\alpha, C, \text{returned}_\alpha(f) \circ Q[\alpha]^f, Q_\ddagger[\alpha]^f \end{array} \right) \\ \wedge (f \in R_F \Rightarrow \forall \alpha' \neq \alpha. c(\alpha') \in \{\text{dead}, \text{halt}\}) \\ \wedge P_\alpha \Rightarrow_\ddagger Q_\ddagger[\alpha]^f \wedge f = \text{last}(\pi_\alpha(s \upharpoonright_{F^\ddagger \& R_F})) \end{array} \right) \right) \quad (\text{IH-Exe})$$

And we want to find a $\rho'_F \in \nu_F$ and for each agent α that is not dead in c' , find a P'_α and prove **(G-Lin)**, **(G-Pa)**, **(G-Idle)**, **(G-Exe)**, **(G-Rec)**, and **(G-Crs)** for (c', Δ', s') , under the condition that $s' \upharpoonright_{F^i \& R_F} \in \mu_F$. By definition of the semantics, it is easy to observe that $s \sqsubseteq s'$, which means $s \upharpoonright_{F^i \& R_F} \sqsubseteq s' \upharpoonright_{F^i \& R_F}$. Moreover, since μ_F is prefix-closed, $s \upharpoonright_{F^i \& R_F} \in \mu_F$ is true and we can freely use above induction hypotheses.

We perform case analysis on the reduction $\langle c, \Delta, s \rangle \xrightarrow{M}_{\nu_E} \langle c', \Delta', s' \rangle$ to prove the inductive step.

* *Invocation*. If the reduction is an invocation to a regular procedure, then there exists $\alpha, f \in F$, and $a \in \text{par}(f)$ such that

$$c[\alpha] = \text{idle} \quad c' = c[\alpha \mapsto M[\alpha]^f] \quad \Delta' = \Delta[\alpha \mapsto [\arg \mapsto a]] \quad s' = s \cdot \alpha:f.$$

We take $\rho'_F = \rho_F \cdot \alpha:f$. By **(IH-Lin)**, we know there exists t_P such that

$$s \upharpoonright_F \cdot t_P \rightsquigarrow_{\dagger F} \rho_F$$

and we prove **(G-Lin)** by

$$s' \upharpoonright_F = s \upharpoonright_F \cdot \alpha:f \cdot t_P \rightsquigarrow_{\dagger F} s \upharpoonright_F \cdot t_P \cdot \alpha:f \rightsquigarrow_{\dagger F} \rho_F \cdot \alpha:f = \rho'_F.$$

We first prove the remaining goals for the α that invokes f . By **(IH-Lin)**, **(IH-Idle)**, and **(IH-Pa)**, there exists a stable P_α such that

$$(c, \Delta, s) \in P_\alpha \subseteq I[\alpha] \subseteq P[\alpha]^f.$$

By definition, we have $(\Delta, s, \rho_F) \text{invoke}_\alpha(f)(\Delta', s', \rho'_F)$, and therefore we define P'_α and prove the left conjunct of **(G-Pa)** by

$$(\Delta', s', \rho'_F) \in \text{invoke}_\alpha(f) \circ P[\alpha]^f \triangleq P'_\alpha.$$

By **(H-FQ)** and lemma **G.5**, we have

$$\begin{aligned} & \text{safe}_\alpha(\mathcal{R}[\alpha], \mathcal{G}[\alpha], P'_\alpha, P'_\alpha, M[\alpha]^f, \text{returned}_\alpha(f) \circ Q[\alpha]^f, Q_\dagger[\alpha]^f) \\ & \text{stable}(\mathcal{R}[\alpha], P'_\alpha) \quad P'_\alpha \Rightarrow_\dagger Q_\dagger[\alpha]^f \end{aligned}$$

which proves **(G-Exe)** and the right conjunct of **(G-Pa)**. The remaining **(G-Idle)**, **(G-Rec)**, and **(G-Crs)** are apparent because the LHS of the implication is false.

For other α' that are not α , we take $P'_{\alpha'} = P_{\alpha'}$ and only need to prove **(G-Lin)**. The remaining **(G-Idle)**, **(G-Exe)**, **(G-Rec)**, and **(G-Crs)** are apparent because truth values of both the LHS and the RHS of implications are not changed compared to those in induction hypotheses. We only need to show

$$(\Delta', s', \rho'_F) \in P_{\alpha'}$$

which is true because $P_{\alpha'}$ is stable (by the left conjunct of **(IH-Pa)**) w.r.t. $\mathcal{R}[\alpha']$, which contains $\text{invoke}_\alpha(f)$ (by **(H-RG)**) that changes (Δ, s, ρ_F) into (Δ', s', ρ'_F) , and $(c, \Delta, s) \in P_{\alpha'}$ is true by **(IH-Pa)**.

* *Return*. If the reduction is a return of a regular procedure, then there exists α, f, v such that

$$c[\alpha] = \text{skip} \quad c' = c[\alpha \mapsto \text{idle}] \quad \Delta(\alpha)(\text{res}) = v \in \text{ar}(f) \quad \Delta' = \Delta[\alpha \mapsto \emptyset]$$

$$\text{last}(\pi_\alpha(s \upharpoonright_{F^i \& R_F})) = f \in F \quad s' = s \cdot \alpha:v.$$

By **(IH-Exe)**, there exists P_α such that

$$\text{safe}_\alpha(\mathcal{R}[\alpha], \mathcal{G}[\alpha], \text{invoke}_\alpha(f) \circ P[\alpha]^f, P_\alpha, \text{skip}, \text{returned}_\alpha(f) \circ Q[\alpha]^f, Q_\dagger[\alpha]^f)$$

By the definition of the safety judgement and **(G-Pa)**, we can further have

$$(\Delta, s, \rho_F) \in P_\alpha \subseteq \text{rely}(\mathcal{R}[\alpha], P_\alpha) \subseteq \text{returned}_\alpha(f) \circ Q[\alpha]^f \circ \text{invoke}_\alpha(f) \circ P[\alpha]^f$$

which means ρ_F already has some $\alpha:v'$ linearized to the end of $\pi_\alpha(\rho_F)$, where $v' = \Delta(\alpha)(\text{res})$ by the definition of returned. Therefore, $\alpha:v$ is linearized since $v' = \Delta(\alpha)(\text{res}) = v$. We take $\rho'_F = \rho_F$ and have

$$(\Delta, s, \rho_F)\text{return}_\alpha(f)(\Delta', s', \rho'_F).$$

By (IH-Lin), there exists t_P such that

$$s \upharpoonright_F \cdot t_P \rightsquigarrow_{\dagger F} \rho_F$$

and $\alpha:v \in t_P$ since $\text{last}(\pi_\alpha(s \upharpoonright_F)) = f$. We take $t'_P = t_P \setminus \alpha:v$ and prove (G-Lin) by

$$s' \upharpoonright_F = s \upharpoonright_F \cdot \alpha:v \cdot t'_P \rightsquigarrow_{\dagger F} s \upharpoonright_F \cdot t_P \rightsquigarrow_{\dagger F} \rho_F = \rho'_F.$$

We first prove remaining goals for the α that returns from f . We define $P'_\alpha = I[\alpha]$ and by our previous argument and (H-FPP), (H-Asrt), we have

$$(\Delta', s', \rho'_F) \in \text{return}_\alpha(f) \circ \text{returned}_\alpha(f) \circ Q[\alpha]^f \circ \text{invoke}_\alpha(f) \circ P[\alpha]^f \subseteq I[\alpha] = P'_\alpha$$

which proves the left conjunct of (G-Pa). By (H-FQ) and lemma G.5, we have

$$\forall f \in F.\text{stable}(\mathcal{R}[\alpha], P[\alpha]^f)$$

which implies $\text{stable}(\mathcal{R}[\alpha], \cap_{f \in F} P[\alpha]^f)$, i.e., $\text{stable}(\mathcal{R}[\alpha], I[\alpha])$. And this proves the right conjunct of (G-Pa).

By (H-Asrt) and reflexivity, we prove (G-Idle). Other branches (G-Crs), (G-Rec), and (G-Exe) are apparent because their LHS of the implication is false.

For other α' that are not α , we take $P'_{\alpha'} = P_{\alpha'}$ and prove (G-Lin), (G-Idle), (G-Exe), (G-Rec), and (G-Crs) with similar argument as the invocation case.

* *Execution.* If the reduction is one step of execution, then there exists α and $C, C' \in \text{Com}$ such that

$$c[\alpha] = C \quad c' = c[\alpha \mapsto C'] \quad \langle C, \Delta, s \rangle \longrightarrow_\alpha \langle C', \Delta', s' \rangle \quad s' \upharpoonright_E \in \nu_E.$$

By unfolding the definition of \longrightarrow_α , we have

$$C \rightsquigarrow_B^X C' \wedge (\Delta', s') \in \llbracket B \rrbracket_\alpha^X(\Delta, s) \quad (\text{H-Cmd})$$

By the definition of $C \rightsquigarrow_B^X C'$, we know $C \neq \text{skip}$. By (IH-Exe), we know there exists some P_α and $f \in F \cup R_F$ such that

$$\text{safe}_\alpha(\mathcal{R}[\alpha]^f, \mathcal{G}[\alpha]^f, \text{invoke}_\alpha(f) \circ P[\alpha]^f, P_\alpha, C, \text{returned}_\alpha(f) \circ Q[\alpha]^f, Q_\ddagger[\alpha]^f).$$

By unfolding the safety judgement's definition and by (H-Cmd), there exists Q' with the followings.

$$\mathcal{G}[\alpha]^f \vdash_\alpha \{\text{rely}(\mathcal{R}[\alpha]^f, P_\alpha)\} B\{Q'\} \quad (\text{H-BT})$$

$$\text{stable}(\mathcal{R}[\alpha]^f, Q' \circ \text{rely}(\mathcal{R}[\alpha]^f, P_\alpha)) \quad (\text{H-MStb})$$

$$Q' \circ \text{rely}(\mathcal{R}[\alpha]^f, P_\alpha) \Rightarrow_\ddagger Q_\ddagger[\alpha]^f \quad (\text{H-Clnto})$$

$$\text{safe}_\alpha(\mathcal{R}[\alpha]^f, \mathcal{G}[\alpha]^f, \text{invoke}_\alpha(f) \circ P[\alpha]^f, Q' \circ \text{rely}(\mathcal{R}[\alpha]^f, P_\alpha), C', \text{returned}_\alpha(f) \circ Q[\alpha]^f, Q_\ddagger[\alpha]^f) \quad (\text{H-C'Safe})$$

By (IH-Pa), we have $(\Delta, s, \rho_F) \in P_\alpha \subseteq \text{rely}(\mathcal{R}[\alpha]^f, P_\alpha)$. By $s' \upharpoonright_E \in \nu_E$ and (H-Cmd), we have $(\Delta', s') \in \llbracket B \rrbracket_\alpha^X(\Delta, s) \cap \nu_E$. And by $s \upharpoonright_{F^\ddagger \& R_F} \in \mu_F$, (H-BT), and the definition of single instruction's judgement, we know there exists ρ' such that

$$(\Delta, s, \rho_F)Q'(\Delta', s', \rho') \quad (\Delta, s, \rho_F)\mathcal{G}[\alpha]^f(\Delta', s', \rho') \quad \rho_F \dashrightarrow \rho'$$

And by **(IH-Lin)** and transitivity of $- \twoheadrightarrow -$, we have $s \upharpoonright_F \twoheadrightarrow \rho_F \twoheadrightarrow \rho'$, i.e., $s \upharpoonright_F \twoheadrightarrow \rho'$. By taking $\rho'_F = \rho'$, we prove **(G-Lin)**.

We take $P'_\alpha = Q' \circ \text{rely}(\mathcal{R}[\alpha]^f, P_\alpha)$. Since $(\Delta, s, \rho_F) \in \text{rely}(\mathcal{R}[\alpha]^f, P_\alpha)$ and $(\Delta, s, \rho_F) Q'(\Delta', s', \rho')$, we know $(\Delta', s', \rho'_F) \in P'_\alpha$. Combined with the stability from **(H-MStb)**, we prove **(G-Pa)** when $f \in F$.

We prove The first and third conjuncts of **(G-Exe)** directly by **(H-C'Safe)** and **(H-CInto)**. The last conjunct of **(G-Exe)** is true by **(IH-Exe)** because the current reduction does not modify $\pi_\alpha(s \upharpoonright_F)$. If $f \in F$, then the second conjunct is true. If $f \in R_F$, because the current reduction does not change other locations in the continuation c , the second conjunct is true by **(IH-Exe)**. This also proves **(G-Pa)** under the condition $f \in R_F$ since the LHS of the implication is false now.

Other branches **(G-Idle)**, **(G-Rec)**, and **(G-Crs)** are apparent because the LHS of the implication is false.

If $f \in F$, for other α' that is not α , we take $P'_{\alpha'} = P_{\alpha'}$ and prove **(G-Lin)**, **(G-Idle)**, **(G-Exe)**, **(G-Rec)**, and **(G-Crs)** with similar argument as previous cases.

If $f \in R_F$, we have shown that any other thread α' are either dead or halt in both c and c' and we need to maintain **(G-Rec)** for them. We take $P'_{\alpha'} = P'_{\alpha'} = Q' \circ \text{rely}(\mathcal{R}[\alpha]^f, P_{\alpha'})$ and **(G-Pa)** is already proved and branches other than **(G-Rec)** are trivially true. We prove **(G-Rec)** by **(H-CInto)**.

* **Crash**. If the reduction is a crash, then

$$\forall \alpha \in s.c'[\alpha] = \text{dead} \quad \forall \alpha \in Y.\alpha \notin s \Rightarrow c'[\alpha] = \text{halt} \quad \Delta' = \Delta_0 \quad s' = s \cdot \downarrow.$$

We take $\rho'_F = \rho_F$. Since $s \upharpoonright_F = s' \upharpoonright_F$ and by **(IH-Lin)**, we know $s' \upharpoonright_F \twoheadrightarrow \rho'_F$ and prove **(G-Lin)**.

For any α , if $c'[\alpha] = \text{dead}$, then we do not need to prove anything for it. If $c'[\alpha] = \text{halt}$, then we know $\alpha \notin s$, which implies $c[\alpha] \in \{\text{idle}, \text{halt}\}$. For these two cases, we only need to prove **(G-Pa)** and **(G-Crs)**. The branches **(G-Idle)**, **(G-Rec)**, and **(G-Exe)** are trivially true because their LHS of the implication is false.

+ $c[\alpha] = \text{idle}$. By **(IH-Pa)** **(IH-Idle)**,

$$(\Delta, s, \rho_F) \in P_\alpha \subseteq I[\alpha].$$

By **(H-PC)**, we have

$$(\Delta, s, \rho_F) \in I[\alpha] \Rightarrow_{\downarrow} Q_{\downarrow}[\alpha].$$

We take $P'_\alpha = Q_{\downarrow}[\alpha]$ and by definition, we have $(\Delta', s', \rho'_F) = (\Delta_0, s \cdot \downarrow, \rho'_F) \in P'_\alpha$ and proves **(G-Pa)**, since the right conjunct of it is trivially true.

By reflexivity, $P'_\alpha = Q_{\downarrow}[\alpha] \subseteq Q_{\downarrow}[\alpha]$, **(G-Crs)** is true.

+ $c[\alpha] = \text{halt}$. It can be further divided into following cases.

- $\forall \alpha'.c[\alpha'] \in \{\text{dead}, \text{halt}\}$, i.e., recovery has not started yet. We take $P'_\alpha = Q_{\downarrow}[\alpha]$. By **(H-RQ)** and lemma **G.5**, we have

$$\forall \alpha. Q_{\downarrow}[\alpha] \Rightarrow_{\downarrow} Q_{\downarrow}[\alpha]$$

Then by **(IH-Pa)**, **(IH-Crs)**, we have

$$(\Delta, s, \rho_F) \in P_\alpha \subseteq Q_{\downarrow}[\alpha] \Rightarrow_{\downarrow} Q_{\downarrow}[\alpha] = P'_\alpha$$

which indicates $(\Delta', s', \rho'_F) = (\Delta_0, s \cdot \downarrow, \rho'_F) \in P'_\alpha$ by definition and proves **(G-Pa)**.

By reflexivity, $P'_\alpha = Q_{\downarrow}[\alpha] \subseteq Q_{\downarrow}[\alpha]$, **(G-Crs)** is true.

– $\exists \alpha', C \in \text{Com.c}[\alpha'] = C$, i.e., recovery has started. By **(IH-Pa)** and **(IH-Rec)**, we know

$$(\Delta, s, \rho_F) \in P_\alpha \Rightarrow_{\dot{z}} Q_{\dot{z}}[\alpha']$$

for some α' . Therefore, we take $P'_\alpha = Q_{\dot{z}}[\alpha']$. And we have $(\Delta', s', \rho'_F) = (\Delta_0, s \cdot \dot{z}, \rho'_F) \in P'_\alpha$ and prove **(G-Pa)**.

By reflexivity, $P'_\alpha = Q_{\dot{z}}[\alpha'] \subseteq Q_{\dot{z}}[\alpha']$, **(G-Crs)** is true.

– $\exists \alpha'. c[\alpha'] = \text{idle}$. This case is impossible, because **idle** and **halt** will not exist in c at the same time by the definition of the semantics.

* Recovery Invocation. If the reduction is an invocation to the recovery, then there exists α such that

$$c[\alpha] = \text{halt} \quad c' = c[\alpha \mapsto M[\alpha]^r] \quad \Delta' = \Delta \quad \text{last}(s) = \dot{z} \quad s' = s \cdot \alpha:r.$$

We take $\rho'_F = \rho_F$ and prove **(G-Lin)** by the same argument as the crash case.

For any α' other than α that are not dead, by lemma G.3, they are in **halt** state. Since the current reduction does not change $c[\alpha']$, $c'[\alpha']$ is still **halt**. And any thread dead are still dead, which showcase the second conjunct in **(G-Exe)** for α . While $c'[\alpha] = M[\alpha]^r$, the LHS of **(G-Idle)**, **(G-Crs)**, and **(G-Exe)** are all false for α' in the new configuration, which means they are proven. And when we found P'_α , we will take $P'_{\alpha'} = P'_\alpha$ and **(G-Pa)** and **(G-Rec)** will be proved for α' when we prove the invariant for α .

We take $P'_\alpha = \text{invoke}_\alpha(r) \circ P_r[\alpha]$. By **(IH-Pa)**, **(IH-Crs)**, **(H-RPre)**, we have

$$(\Delta, s, \rho_F) \in P_\alpha \subseteq Q_{\dot{z}}[\alpha] \subseteq P_r[\alpha].$$

And by definition, we have

$$(\Delta, s, \rho_F) \text{invoke}_\alpha(r) (\Delta', s', \rho'_F)$$

which implies $(\Delta', s', \rho'_F) \in \text{invoke}_\alpha(r) \circ P_r[\alpha] = P'_\alpha$ and proves **(G-Pa)** for both α and α' .

By **(H-RQ)** and lemma G.5, we have

$$\begin{aligned} & \text{safe}_\alpha(\mathcal{R}[\alpha]^r, \mathcal{G}[\alpha]^r, P'_\alpha, P'_\alpha, M[\alpha]^r, \text{returned}_\alpha(r) \circ Q[\alpha]^r, Q_{\dot{z}}[\alpha]^r) \\ & \text{stable}(\mathcal{R}[\alpha]^r, P'_\alpha) \quad P'_\alpha \Rightarrow_{\dot{z}} Q_{\dot{z}}[\alpha]^r \end{aligned}$$

which proves **(G-Rec)** for α' , and proves **(G-Exe)** for α (the second conjunct is proved in last paragraph). Other branches **(G-Idle)** and **(G-Crs)** are trivially true.

* Recovery Return. If the reduction is a return from the recovery, then there exists α and v such that

$$c[\alpha] = \text{skip} \quad \forall \alpha \in \Upsilon. c[\alpha] = \text{dead} \Rightarrow c'[\alpha] = \text{dead} \quad \forall \alpha \in \Upsilon. c[\alpha] \neq \text{dead} \Rightarrow c'[\alpha] = \text{idle}$$

$$\Delta(\alpha)(\text{res}) = v \in \text{ar}(r) \quad \Delta' = \Delta[\alpha \mapsto \emptyset] \quad \text{last}(\pi_\alpha(s \upharpoonright_{F \cup R_F})) = r \quad s' = s \cdot \alpha:v.$$

We take $\rho'_F = \rho_F$ and prove **(G-Lin)** by the same argument as the crash case.

We take $P'_{\alpha'} = I[\alpha']$ for any α' that are not dead.

By **(IH-Exe)**, there exists P_α such that

$$\text{safe}_\alpha(\mathcal{R}[\alpha]^r, \mathcal{G}[\alpha]^r, \text{invoke}_\alpha(r) \circ P[\alpha]^r, P_\alpha, \text{skip}, \text{returned}_\alpha(r) \circ Q[\alpha]^r, Q_{\dot{z}}[\alpha]^r)$$

By the definition of the safety judgement and **(IH-Pa)**, we can further have

$$(\Delta, s, \rho_F) \in P_\alpha \subseteq \text{rely}(\mathcal{R}[\alpha]^r, P_\alpha) \subseteq \text{returned}_\alpha(r) \circ Q[\alpha]^r \circ \text{invoke}_\alpha(r) \circ P[\alpha]^r$$

By definition, we have

$$(\Delta, s, \rho_F)\text{return}_\alpha(f)(\Delta', s', \rho'_F).$$

By (H-RPost), we have the following for any α' .

$$(\Delta', s', \rho'_F) \in \text{return}_\alpha(f) \circ \text{returned}_\alpha(r) \circ Q[\alpha]^r \circ \text{invoke}_\alpha(r) \circ P[\alpha]^r \subseteq I[\alpha'] = P'_\alpha$$

which proves (G-Pa) for any α' .

Moreover, since any non-dead α' is in idle state, we only need to prove (G-Idle), which is trivially true by reflexivity, $P'_\alpha = I[\alpha'] \subseteq I[\alpha']$.

This concludes the proof of proposition G.7, which proves the auxiliary soundness. □

THEOREM G.8 (SOUNDNESS). *Proposition G.1 is true, i.e., the program logic is sound.*

PROOF. We prove the soundness by the following derivation.

$$\begin{aligned} v'_F \uparrow_{F^\sharp} &= (v'_E; \llbracket M \rrbracket_{R_E} \cap \mu_F) \uparrow_{F^\sharp} && \text{By Definition of } v'_F \\ &= ((v'_E; \llbracket M \rrbracket_{R_E}) \uparrow_{E^\sharp} \multimap_{F \cup R_F} \cap \mu_F) \uparrow_{F^\sharp} \\ &\subseteq (v'_E \uparrow_{E^\sharp}; \llbracket M \rrbracket_\emptyset \cap \mu_F) \uparrow_{F^\sharp} && (1) \text{ Recovery Refinement} \\ &\subseteq (v_E; \llbracket M \rrbracket_\emptyset \cap \mu_F) \uparrow_{F^\sharp} && (2) \text{ Validity of } (v'_E, v_E) + \text{Obs. Ref.} \\ &= (\llbracket \text{Link } v_E; M \rrbracket \cap \mu_F) \uparrow_{F^\sharp} && (3) \text{ Linking (lemma G.2)} \\ &\subseteq \text{dur}(v_F) && (4) \text{ Auxiliary Soundness (lemma G.6)} \end{aligned}$$

□

H A PROGRAM LOGIC FOR CRASH-AWARE OVERLAY OBJECTS

In the main text, we mainly present the program logic for durable overlay objects. However, it is often necessary to verify crash-aware objects, e.g., volatile objects and buffered objects, which are crucial for implementations of many objects. Our file system example is a two layer crash-aware object implementation. Therefore, we propose a program logic for crash-aware linearizability and crash-aware overlay objects in this section.

We use the same programming language and semantics from G for crash-aware objects. For simplicity, we still enforce the durable assumption here so that we can reuse some previous definitions and conclusions.

H.1 Interfaces

We use the same interface definitions for crash-aware objects in G, which we repeat here. The interface of a crash-aware linearizable object E is a tuple

$$(v'_E : \dagger(E \cup R_E) \in \mathbf{Dur}, v_E : \dagger E^\sharp \in \mathbf{Crash}) \quad \text{s.t.} \quad v'_E \uparrow_{E^\sharp} \subseteq K_d v_E$$

where v'_E is the concrete specification that contains all possible traces the object can produce, which will include concurrent ones and will also contain crash events and recovery signatures, and v_E is the linearized specification. The interface is valid if and only if after recovery refining the concrete specification (the projection onto E^\sharp), $v'_E \uparrow_{E^\sharp}$ is crash-aware linearizable to v_E , i.e., $v'_E \uparrow_{E^\sharp} \subseteq K_d v_E$.

The objective for the program logic is to establish the judgement:

$$\mu_F \vdash^{\text{ca}} M : (v'_E, v_E) \rightarrow (v'_F, v_F)$$

where μ_F is the client guarantee. The soundness will ensure (v'_F, v_F) to be valid given a valid underlay (v'_E, v_E) .

H.2 The CLHL for Crash-Aware Linearizability

The program logic uses as proof configurations triples $(\Delta, s, \rho) \in \text{Config} := \text{ModState} \times \text{Poss}$, where Poss is a set of possibilities and is of type $\dagger F^{\frac{1}{2}}$. The possibility is different from the one in the program logic for durable objects, where we need to consider crash events in the linearization. We define assertions and rely-guarantees as sets and relations just like before.

The program logic is almost the same as the one for the program logic for durable objects, except one crucial difference: **we are using a different rewrite system now.**

Definition H.1. Let $A = (M_A, P_A)$ be a crash-aware game. We define a string rewrite system $(P_A, \rightsquigarrow_A^{\frac{1}{2}})$ with rewrite rules:

- $\forall m, m' \in M_A. \forall \alpha, \alpha' \in \Upsilon. \forall X \in \{O, P\}. \alpha \neq \alpha' \wedge \lambda_A(m) = \alpha : X \wedge \lambda_A(m') = \alpha' : X \implies m \cdot m' \rightsquigarrow_A^{\frac{1}{2}} m' \cdot m$
- $\forall o, p \in M_A. \forall \alpha, \alpha' \in \Upsilon. \alpha \neq \alpha' \wedge \lambda_A(o) = \alpha : O \wedge \lambda_A(p) = \alpha' : P \implies o \cdot p \rightsquigarrow_A^{\frac{1}{2}} p \cdot o$
- $\forall m \in M_A^Y. \forall \Upsilon \in \Upsilon. \lambda_A(m) = \alpha : P \implies \frac{1}{2} \cdot m \rightsquigarrow_A^{\frac{1}{2}} m \cdot \frac{1}{2}$

It differs from the rewrite system for concurrent games mainly in the third rule, which allows insertion of a proponent action before a crash, if this insertion is valid. We define the crash-aware ghost update as

$$s \dashrightarrow^{\frac{1}{2}} t \iff \exists s_P \in (M_F^P)^*. s \cdot s_P \rightsquigarrow_{\dagger F^{\frac{1}{2}}}^{\frac{1}{2}} t$$

and lemma H.2 establishes the connection between the crash-aware update and crash-aware linearizability.

LEMMA H.2. *For any crash-aware play s and t in $\dagger F^{\frac{1}{2}}$, $s \dashrightarrow^{\frac{1}{2}} t \implies s \rightsquigarrow^{\frac{1}{2}} t$.*

For the configuration triple (Δ, s, ρ) , we maintain the invariant that $s \dashrightarrow^{\frac{1}{2}} \rho$, i.e., $s \rightsquigarrow^{\frac{1}{2}} \rho$, and $\rho \in \nu_F$.

There are two differences in the program logic's proof rules. One is the primitive judgement $\mathcal{G} \vdash_{\alpha}^{\text{ca}} \{P\}B\{Q\}$, where $B \in \text{Prim}$. We now use the crash-aware ghost update to linearize the possibility.

$$\mathcal{G} \vdash_{\alpha}^{\text{ca}} \{P\}B\{Q\} \iff \left(\begin{array}{l} \forall (\Delta, s, \rho). s \upharpoonright_{F \cup R_F} \in \mu_F \wedge (\Delta, s, \rho) \in P \wedge \\ \forall (\Delta', s') \in \llbracket B \rrbracket_{\alpha}(\Delta, s) \cap \nu_E \implies s' \upharpoonright_{F \cup R_F} \in \mu_F \wedge \\ \exists \rho'. (\Delta, s, \rho)Q(\Delta', s', \rho') \wedge (\Delta, s, \rho)\mathcal{G}(\Delta', s', \rho') \wedge \rho \dashrightarrow^{\frac{1}{2}} \rho' \end{array} \right)$$

Another one is the crash-into relation $(- \implies_{\frac{1}{2}}^{\text{ca}} -)$, where we also add the crash event to the possibility as well. And we replace with this crash-into relation in **PRIM** rule, **SKIP** rule, **CONSEQ** rule, and **OBJECT IMPL** rule, obtaining the program logic for crash-aware linearizability.

$$P \implies_{\frac{1}{2}}^{\text{ca}} Q_{\frac{1}{2}} \iff \forall (\Delta, s, \rho) \in P. (\Delta_0, s \cdot \frac{1}{2}, \rho \cdot \frac{1}{2}) \in Q_{\frac{1}{2}}$$

Essential, the usage of this program logic is the same as the one for durable objects, except we use different possibility update operation when establishing the primitive judgement and need to prove with a slightly different crash-into relation. Fig. 13 shows all proof rules in the program logic for crash-aware objects.

Crash Hoare Logic Rules:

$$\begin{array}{c}
\frac{P \Rightarrow_{\frac{ca}{\downarrow}}^{\text{ca}} Q_{\downarrow} \quad Q \circ P \Rightarrow_{\frac{ca}{\downarrow}}^{\text{ca}} Q_{\downarrow} \quad Q_{\downarrow} \Rightarrow_{\frac{ca}{\downarrow}}^{\text{ca}} Q_{\downarrow}}{\text{stable}(\mathcal{R}, P) \quad \text{stable}(\mathcal{R}, Q) \quad \mathcal{G} \vdash_{\alpha}^{\text{ca}} \{P\}B\{Q\}}}{\mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}B\{Q\}\{Q_{\downarrow}\}} \text{ PRIM} \\
\\
\frac{\mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}C_1\{Q_1\}\{Q_{\downarrow}\} \quad \mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{Q_1 \circ P\}C_2\{Q_2\}\{Q_{\downarrow}\}}{\mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}C_1; C_2\{Q_2 \circ Q_1\}\{Q_{\downarrow}\}} \text{ SEQ} \\
\\
\frac{\text{stable}(\mathcal{R}, P) \quad P \Rightarrow_{\frac{ca}{\downarrow}}^{\text{ca}} Q_{\downarrow}}{\mathcal{R}, \text{ID} \vDash_{\alpha}^{\text{ca}} \{P\}\text{skip}\{\text{ID}\}\{Q_{\downarrow}\}} \text{ SKIP} \quad \frac{\mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}C\{Q\}\{Q_{\downarrow}\} \quad Q \circ P \subseteq P}{\mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}C^*\{Q\}\{Q_{\downarrow}\}} \text{ ITER} \\
\\
\frac{\mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}C_1\{Q\}\{Q_{\downarrow}\} \quad \mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}C_2\{Q\}\{Q_{\downarrow}\}}{\mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}C_1 + C_2\{Q\}\{Q_{\downarrow}\}} \text{ CHOICE} \\
\\
\frac{\text{stable}(\mathcal{R}', P') \quad \text{stable}(\mathcal{R}', Q') \quad Q_{\downarrow} \subseteq Q'_{\downarrow} \quad Q'_{\downarrow} \Rightarrow_{\frac{ca}{\downarrow}}^{\text{ca}} Q'_{\downarrow} \quad Q' \circ P' \Rightarrow_{\frac{ca}{\downarrow}}^{\text{ca}} Q'_{\downarrow}}{P' \subseteq P \quad Q \subseteq Q' \quad \mathcal{R}' \subseteq \mathcal{R} \quad \mathcal{G}' \subseteq \mathcal{G}' \quad \mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}C\{Q\}\{Q_{\downarrow}\}}{\mathcal{R}', \mathcal{G}' \vDash_{\alpha}^{\text{ca}} \{P'\}C\{Q'\}\{Q'_{\downarrow}\}} \text{ CONSEQ}
\end{array}$$

Top Level Rules:

$$\begin{array}{c}
\frac{\forall f \in F. (\Delta_0, \epsilon, \epsilon) \in P[\alpha]^f \quad \forall f \in F. P[\alpha]^f \subseteq \text{idle}_{\alpha} \quad \forall f \in F. \text{stable}(\mathcal{R}[\alpha], P[\alpha]^f) \\
\forall f \in F. \mathcal{R}[\alpha], \mathcal{G}[\alpha] \vDash_{\alpha}^{\text{ca}} \{\text{invoke}_{\alpha}(f) \circ P[\alpha]^f\}M[\alpha]^f\{\text{returned}_{\alpha}(f) \circ Q[\alpha]^f\}\{\top\} \\
\forall f, f' \in F. \text{return}_{\alpha}(f) \circ \text{returned}_{\alpha}(f) \circ Q[\alpha]^f \circ \text{invoke}_{\alpha}(f) \circ P[\alpha]^f \subseteq P[\alpha]^{f'}}}{\mathcal{R}[\alpha], \mathcal{G}[\alpha], (\cap_{f \in F} P[\alpha]^f) \vDash_{\alpha}^{\text{ca}} M_F[\alpha]} \text{ LOCAL IMPL} \\
\\
\frac{\forall \alpha \in \Upsilon. \mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \vDash_{\alpha}^{\text{ca}} M_F[\alpha] \\
\forall \alpha, \alpha' \in \Upsilon. \alpha \neq \alpha' \Rightarrow \mathcal{G}[\alpha] \cup \text{invoke}_{\alpha}(-) \cup \text{return}_{\alpha}(-) \subseteq \mathcal{R}[\alpha] \\
\forall \alpha. I[\alpha] \Rightarrow_{\frac{ca}{\downarrow}}^{\text{ca}} Q_{\downarrow}[\alpha] \quad \forall \alpha, \alpha' \in \Upsilon. Q_{\downarrow}[\alpha] \subseteq P_r[\alpha'] \\
\forall \alpha \in \Upsilon. \text{ID}, \top \vDash_{\alpha}^{\text{ca}} \{\text{invoke}_{\alpha}(r) \circ P_r[\alpha]\}M[\alpha]^r\{\text{returned}_{\alpha}(r) \circ Q_r[\alpha]\}\{Q_{\downarrow}[\alpha]\} \\
\forall \alpha, \alpha' \in \Upsilon. \text{return}_{\alpha}(r) \circ \text{returned}_{\alpha}(r) \circ Q_r[\alpha] \circ \text{invoke}_{\alpha}(r) \circ P_r[\alpha] \subseteq I[\alpha']}{\mu F \vdash^{\text{ca}} M : (v'_E, v_d, v_c) \rightarrow (v'_F, v_F)} \text{ OBJECT IMPL}
\end{array}$$

Fig. 13. Proof Rules in the Program Logic for Crash-Aware Objects

H.3 Soundness

The program logic is justified by the following soundness theorem.

PROPOSITION H.3 (SOUNDNESS). *If $\mu F \vdash^{\text{ca}} M : (v'_E, v_E) \rightarrow (v'_F, v_F)$ is provable, and (v'_E, v_E) is a valid underlay interface, and $v'_F = v'_E; \llbracket M \rrbracket_{R_E} \cap \mu F$, then (v'_F, v_F) is a valid overlay interface with*

$$v'_F \upharpoonright_{F^{\downarrow}} \subseteq K_{\frac{ca}{\downarrow}} v_F.$$

The soundness proof of the crash-aware program logic is also almost identical to the one of the durable program logic in **G**. We use the same auxiliary definitions for judgement of the crash-aware version and the only difference is the auxiliary soundness (lemma **H.4**), because it is the only one mentions the only thing different between the two logics, the linearization (i.e., the possibility update).

$$\begin{aligned}
v'_F \upharpoonright_{F^\ddagger} &= (v'_E; \llbracket M \rrbracket_{R_E} \cap \mu_F) \upharpoonright_{F^\ddagger} && \text{By Definition of } v'_F \\
&= ((v'_E; \llbracket M \rrbracket_{R_E}) \upharpoonright_{E^\ddagger} \multimap_{F \cup R_F} \cap \mu_F) \upharpoonright_{F^\ddagger} \\
&\subseteq (v'_E \upharpoonright_{E^\ddagger}; \llbracket M \rrbracket_\emptyset \cap \mu_F) \upharpoonright_{F^\ddagger} && (1) \text{ Recovery Refinement} \\
&\subseteq (v_E; \llbracket M \rrbracket_\emptyset \cap \mu_F) \upharpoonright_{F^\ddagger} && (2) \text{ Validity of } (v'_E, v_E) + \text{Obs. Ref.} \\
&= (\llbracket \text{Link } v_E; M \rrbracket \cap \mu_F) \upharpoonright_{F^\ddagger} && (3) \text{ Linking (lemma G.2)} \\
&\subseteq K_\ddagger v_F && (4) \text{ Auxiliary Soundness (lemma H.4)}
\end{aligned}$$

LEMMA H.4 (AUXILIARY SOUNDNESS). *If the judgement $\mu_F \vdash^{\text{ca}} M : (v'_E, v_E) \rightarrow (v'_F, v_F)$ is provable and (v'_E, v_E) is a valid underlay interface, then*

$$(\llbracket \text{Link } v_E; M \rrbracket \cap \mu_F) \upharpoonright_{F^\ddagger} \subseteq K_\ddagger v_F.$$

PROOF. The proof of this auxiliary soundness is almost the same as the proof of lemma G.6 for the durable program logic. We maintain the invariant that for any c, Δ, s , if $\langle c_0, \Delta_0, \epsilon \rangle \xrightarrow{M_{v_E}} \langle c, \Delta, s \rangle$, then when $s \upharpoonright_{F^\ddagger \& R_F} \in \mu_F$ there exists a current linearization $\rho_F \in v_F$ and the followings hold.

$$\begin{aligned}
s \upharpoonright_{F^\ddagger} &\xrightarrow{\ddagger} \rho_F && \text{(G-Lin)} \\
\bigwedge \forall \alpha. c[\alpha] &\neq \text{dead} \Rightarrow && \\
\exists P_\alpha. (\Delta, s, \rho_F) &\in P_\alpha \wedge (\text{halt} \notin c \Rightarrow \text{stable}(\mathcal{R}[\alpha], P_\alpha)) && \text{(G-Pa)} \\
&\wedge (c[\alpha] = \text{idle} \Rightarrow P_\alpha \subseteq P[\alpha]) && \text{(G-Idle)} \\
&\wedge ((\forall \alpha'. c[\alpha'] \in \{\text{dead}, \text{halt}\}) \Rightarrow \exists \alpha'. P_\alpha \subseteq Q_\ddagger[\alpha']) && \text{(G-Crs)} \\
&\wedge (c[\alpha] = \text{halt} \wedge (\exists \alpha', C \in \text{Com}.c[\alpha'] = C) \Rightarrow \exists \alpha'. P_\alpha \Rightarrow_{\ddagger}^{\text{ca}} Q_\ddagger[\alpha']) && \text{(G-Rec)} \\
&\wedge \left(\forall C \in \text{Com}.c[\alpha] = C \Rightarrow \exists f \in F \cup R_F. \left(\begin{aligned} &\text{safe}_\alpha \left(\mathcal{R}[\alpha]^f, \mathcal{G}[\alpha]^f, \text{invoke}_\alpha(f) \circ P[\alpha]^f, \right) \\ &P_\alpha, C, \text{returned}_\alpha(f) \circ Q[\alpha]^f, Q_\ddagger[\alpha]^f \right) \\ &\wedge (f \in R_F \Rightarrow \forall \alpha' \neq \alpha. c[\alpha'] \in \{\text{dead}, \text{halt}\}) \\ &\wedge P_\alpha \Rightarrow_{\ddagger}^{\text{ca}} Q_\ddagger[\alpha]^f \wedge f = \text{last}(\pi_\alpha(s \upharpoonright_{F^\ddagger \& R_F})) \end{aligned} \right) \right) && \text{(G-Exe)}
\end{aligned}$$

Combined with lemma H.2 and the crash-aware linearizability's definition in §B, this invariant proves the auxiliary soundness by showing

$$\forall s \in (\llbracket \text{Link } v_E; M \rrbracket \cap \mu_F). s \upharpoonright_{F^\ddagger} \in K_\ddagger v_F.$$

We prove the invariant by first inducting over the reduction $\langle c_0, \Delta_0, \epsilon \rangle \xrightarrow{M_{v_E}^*} \langle c, \Delta, s \rangle$ and in the inductive step where we have

$$\langle c_0, \Delta_0, \epsilon \rangle \xrightarrow{M_{v_E}^*} \langle c, \Delta, s \rangle \text{ and } \langle c, \Delta, s \rangle \xrightarrow{M_{v_E}^*} \langle c', \Delta', s' \rangle.$$

We then prove by a case analysis of the last reduction step. The major difference between this proof and the one in G is the case where the last reduction step is a crash. We only demonstrate the proof for the crash case here.

We have the induction hypothesis that there exists a current linearization $\rho_F \in \nu_F$ and the followings hold.

$$\begin{aligned}
& s \upharpoonright_{F^i} \xrightarrow{\not\downarrow} \rho_F && \text{(IH-Lin)} \\
& \wedge \forall \alpha. c[\alpha] \neq \text{dead} \Rightarrow \\
& \quad \exists P_{\alpha}. (\Delta, s, \rho_F) \in P_{\alpha} \wedge (\text{halt} \notin c \Rightarrow \text{stable}(\mathcal{R}[\alpha], P_{\alpha})) && \text{(IH-Pa)} \\
& \quad \wedge (c[\alpha] = \text{idle} \Rightarrow P_{\alpha} \subseteq P[\alpha]) && \text{(IH-Idle)} \\
& \quad \wedge ((\forall \alpha'. c[\alpha'] \in \{\text{dead}, \text{halt}\}) \Rightarrow \exists \alpha'. P_{\alpha} \subseteq Q_{\not\downarrow}[\alpha']) && \text{(IH-Crs)} \\
& \quad \wedge (c[\alpha] = \text{halt} \wedge (\exists \alpha', C \in \text{Com}.c[\alpha'] = C) \Rightarrow \exists \alpha'. P_{\alpha} \Rightarrow_{\not\downarrow}^{\text{ca}} Q_{\not\downarrow}[\alpha']) && \text{(IH-Rec)} \\
& \quad \wedge \left(\forall C \in \text{Com}.c[\alpha] = C \Rightarrow \exists f \in F \cup R_F. \left(\begin{array}{l} \text{safe}_{\alpha} \left(\begin{array}{l} \mathcal{R}[\alpha]^f, \mathcal{G}[\alpha]^f, \text{invoke}_{\alpha}(f) \circ P[\alpha]^f, \\ P_{\alpha}, C, \text{returned}_{\alpha}(f) \circ Q[\alpha]^f, Q_{\not\downarrow}[\alpha]^f \end{array} \right) \\ \wedge (f \in R_F \Rightarrow \forall \alpha' \neq \alpha. c[\alpha'] \in \{\text{dead}, \text{halt}\}) \\ \wedge P_{\alpha} \Rightarrow_{\not\downarrow}^{\text{ca}} Q_{\not\downarrow}[\alpha]^f \wedge f = \text{last}(\pi_{\alpha}(s \upharpoonright_{F^i \& R_F})) \end{array} \right) \right) && \text{(IH-Exe)}
\end{aligned}$$

We also list some premise extracted from the judgement $\mu_F \vdash^{\text{ca}} M : (\nu'_E, \nu_E) \rightarrow (\nu'_F, \nu_F)$ that are used in the crash case's proof.

$$\begin{aligned}
& \forall \alpha \in A.\text{ID}, \top \vDash_{\alpha}^{\text{ca}} \{\text{invoke}_{\alpha}(r) \circ P_r[\alpha]\} M[\alpha]^r \{\text{returned}_{\alpha}(r) \circ Q_r[\alpha]\} \{Q_{\not\downarrow}[\alpha]\} && \text{(H-RQ)} \\
& \quad \forall \alpha. I[\alpha] \Rightarrow_{\not\downarrow}^{\text{ca}} Q_{\not\downarrow}[\alpha] && \text{(H-PC)}
\end{aligned}$$

We also have the crash-aware version of lemma G.5 as lemma H.5

LEMMA H.5. *For any $\mathcal{R}, \mathcal{G}, P, s, Q, Q_{\not\downarrow}$, if the quadruple $\mathcal{R}, \mathcal{G} \vDash_{\alpha}^{\text{ca}} \{P\}s\{Q\}\{Q_{\not\downarrow}\}$ is provable, then the followings are true.*

$$\text{stable}(\mathcal{R}, P) \quad P \Rightarrow_{\not\downarrow}^{\text{ca}} Q_{\not\downarrow} \quad Q_{\not\downarrow} \Rightarrow_{\not\downarrow}^{\text{ca}} Q_{\not\downarrow} \quad \text{safe}_{\alpha}^{\text{ca}}(\mathcal{R}, \mathcal{G}, P, P, s, Q, Q_{\not\downarrow})$$

* Crash. If the reduction is a crash, then

$$\forall \alpha \in s.c'[\alpha] = \text{dead} \quad \forall \alpha \in \Upsilon.\alpha \notin s \Rightarrow c'[\alpha] = \text{halt} \quad \Delta' = \Delta_0 \quad s' = s \cdot \not\downarrow.$$

We have the induction hypothesis that there exists $\rho_F \in \nu_F$, such that

$$s \upharpoonright_{F^i} \xrightarrow{\not\downarrow} \rho_F$$

We take $\rho'_F = \rho_F \cdot \not\downarrow$. By definition, the induction hypothesis is equivalent to

$$\exists t_P \in (M_F^P)^*. s \upharpoonright_{F^i} \cdot t_P \rightsquigarrow_{\not\downarrow}^{\not\downarrow} \rho_F.$$

By the third rewrite rule of $- \rightsquigarrow_{\not\downarrow}^{\not\downarrow} -$, we insert t_P from previous linearization before the crash and have

$$s' \upharpoonright_{F^i} = s \upharpoonright_{F^i} \cdot \not\downarrow \rightsquigarrow_{\not\downarrow}^{\not\downarrow} s \upharpoonright_{F^i} \cdot t_P \cdot \not\downarrow \rightsquigarrow_{\not\downarrow}^{\not\downarrow} \rho_F \cdot \not\downarrow = \rho'_F$$

which shows $s' \upharpoonright_{F^i} \xrightarrow{\not\downarrow} \rho'_F$ by definition and proves (G-Lin).

For any α , if $c'[\alpha] = \text{dead}$, then we do not need to prove anything for it. If $c'[\alpha] = \text{halt}$, then we know $\alpha \notin s$, which implies $c[\alpha] \in \{\text{idle}, \text{halt}\}$. For these two cases, we only need to prove (G-Pa) and (G-Crs). The branches (G-Idle), (G-Rec), and (G-Exe) are trivially true because their LHS of the implication is false.

+ $c[\alpha] = \text{idle}$. By (IH-Pa) (IH-Idle),

$$(\Delta, s, \rho_F) \in P_\alpha \subseteq I[\alpha].$$

By (H-PC), we have

$$(\Delta, s, \rho_F) \in I[\alpha] \Rightarrow_{\downarrow}^{\text{ca}} Q_{\downarrow}[\alpha].$$

We take $P'_\alpha = Q_{\downarrow}[\alpha]$ and by definition, we have $(\Delta', s', \rho'_F) = (\Delta_0, s \cdot \downarrow, \rho_F \cdot \downarrow) \in P'_\alpha$ and proves (G-Pa), since the right conjunct of it is trivially true.

By reflexivity, $P'_\alpha = Q_{\downarrow}[\alpha] \subseteq Q_{\downarrow}[\alpha]$, (G-Crs) is true.

+ $c[\alpha] = \text{halt}$. It can be further divided into following cases.

- $\forall \alpha'. c[\alpha'] \in \{\text{dead}, \text{halt}\}$, i.e., recovery has not started yet. We take $P'_\alpha = Q_{\downarrow}[\alpha]$. By (H-RQ) and lemma H.5, we have

$$\forall \alpha. Q_{\downarrow}[\alpha] \Rightarrow_{\downarrow}^{\text{ca}} Q_{\downarrow}[\alpha]$$

Then by (IH-Pa), (IH-Crs), we have

$$(\Delta, s, \rho_F) \in P_\alpha \subseteq Q_{\downarrow}[\alpha] \Rightarrow_{\downarrow}^{\text{ca}} Q_{\downarrow}[\alpha] = P'_\alpha$$

which indicates $(\Delta', s', \rho'_F) = (\Delta_0, s \cdot \downarrow, \rho_F \cdot \downarrow) \in P'_\alpha$ by definition and proves (G-Pa).

By reflexivity, $P'_\alpha = Q_{\downarrow}[\alpha] \subseteq Q_{\downarrow}[\alpha]$, (G-Crs) is true.

- $\exists \alpha', C \in \text{Com}. c[\alpha'] = C$, i.e., recovery has started. By (IH-Pa) and (IH-Rec), we know

$$(\Delta, s, \rho_F) \in P_\alpha \Rightarrow_{\downarrow}^{\text{ca}} Q_{\downarrow}[\alpha']$$

for some α' . Therefore, we take $P'_\alpha = Q_{\downarrow}[\alpha']$. And we have

$$(\Delta', s', \rho'_F) = (\Delta_0, s \cdot \downarrow, \rho_F \cdot \downarrow) \in P'_\alpha$$

and prove (G-Pa). By reflexivity, $P'_\alpha = Q_{\downarrow}[\alpha'] \subseteq Q_{\downarrow}[\alpha']$, (G-Crs) is true.

- $\exists \alpha'. c[\alpha'] = \text{idle}$. This case is impossible, because idle and halt will not exist in c at the same time by the definition of the semantics.

□

I APPLICATIONS OF THE PROGRAM LOGIC

In this section, we demonstrate the ability of our program logic from appendix G by presenting detailed proofs of the FLiT example and the file system example in §4.2. In I.1, we formalize and verify the FLiT memory cell implementation. And in I.3, we formalize and verify the file system example.

I.1 The FLiT Memory Cell

I.1.1 Implementation. The FLiT memory cell implementation has the signature:

$$\text{FLiT} := \{\text{load} : \text{Val}, \text{store} : \text{Val} \rightarrow \mathbf{1}\}$$

It is durably linearizable to its specification v_{flit} , which is the largest set satisfying the following property.

$$p \in v_{\text{flit}} \iff \left(\begin{array}{c} p = \epsilon \vee \\ (p \sqsubseteq p' \cdot \alpha : \text{store}(v) \cdot \alpha : \text{ok} \wedge p' \in v_{\text{flit}}) \vee \\ (p \sqsubseteq p' \cdot \alpha : \text{load} \cdot \alpha : v \wedge p' \in v_{\text{flit}} \wedge v = \text{fstate}(p')) \end{array} \right)$$

$$\text{where } \text{fstate}(p) = \begin{cases} v_0 & p = \epsilon \\ v & p = p' \cdot \alpha : \text{ok} \wedge \pi_\alpha(p') = p'' \cdot \text{store}(v) \\ \text{fstate}(p') & \text{otherwise, } p = p' \cdot _ \end{cases}$$

It builds on a volatile and atomic counter Counter and a buffered and atomic memory cell BCell with the following signatures:

$$\text{Counter} := \{\text{inc} : \mathbf{1}, \text{dec} : \mathbf{1}, \text{get} : \mathbb{Z}\}$$

$$\text{BCell} := \{\text{load} : \text{Val}, \text{store} : \text{Val} \rightarrow \mathbf{1}, \text{flush} : \mathbf{1}\}$$

The counter has the specification v_{counter} as the largest set satisfying the following property.

$$s \in v_{\text{counter}} \iff \left(\begin{array}{c} s = \epsilon \vee (s = s' \cdot \frac{1}{2} \wedge s' \in v_{\text{counter}}) \vee \\ (s \sqsubseteq s' \cdot \alpha : \text{inc} \cdot \alpha : \text{ok} \wedge s' \in v_{\text{counter}}) \vee \\ (s \sqsubseteq s' \cdot \alpha : \text{dec} \cdot \alpha : \text{ok} \wedge s' \in v_{\text{counter}}) \vee \\ (s \sqsubseteq s' \cdot \alpha : \text{get} \cdot \alpha : n \wedge s' \in v_{\text{counter}} \wedge n = \text{cstate}(s')) \end{array} \right)$$

$$\text{where } \text{cstate}(s) = \begin{cases} 0 & s = \epsilon \vee s = s' \cdot \frac{1}{2} \\ \text{cstate}(s') + 1 & s = s' \cdot \alpha : \text{ok} \wedge \pi_\alpha(s') = s'' \cdot \text{inc} \\ \text{cstate}(s') - 1 & s = s' \cdot \alpha : \text{ok} \wedge \pi_\alpha(s') = s'' \cdot \text{dec} \\ \text{cstate}(s') & \text{otherwise, } s = s' \cdot _ \end{cases}$$

The buffered memory cell has the specification v_{bcell} as the largest set satisfying the following property.

$$s \in v_{\text{bcell}} \iff \left(\begin{array}{c} s = \epsilon \vee (s = s' \cdot \frac{1}{2} \wedge s' \in v_{\text{bcell}}) \vee \\ (s \sqsubseteq s' \cdot \alpha : \text{store}(v) \cdot \alpha : \text{ok} \wedge s' \in v_{\text{bcell}}) \vee \\ (s \sqsubseteq s' \cdot \alpha : \text{flush} \cdot \alpha : \text{ok} \wedge s' \in v_{\text{bcell}}) \vee \\ (s \sqsubseteq s' \cdot \alpha : \text{load} \cdot \alpha : v \wedge s' \in v_{\text{bcell}} \wedge v \in \text{mstate}(s') \upharpoonright_2) \end{array} \right)$$

$$\text{where } (v_1, v_2) \in \text{mstate}(s) \iff \left(\begin{array}{c} (v_1 = v_2 = v_0 \wedge s = \epsilon) \vee \\ (v_1 = v_2 = v \wedge s = s' \cdot \frac{1}{2} \wedge (v, v') \in \text{mstate}(s')) \vee \\ \left(\begin{array}{c} v_1 = v' \wedge v_2 = v \wedge s = s' \cdot \alpha : \text{ok} \wedge \\ \pi_\alpha(s') = s'' \cdot \text{store}(v) \wedge (v', v'') \in \text{mstate}(s') \end{array} \right) \vee \\ (v_1 = v_2 = v \wedge s = s' \cdot \text{ok} \wedge \pi_\alpha(s') = s'' \cdot \text{store}(v)) \vee \\ \left(\begin{array}{c} v_1 = v_2 = v \wedge s = s' \cdot \alpha : \text{ok} \wedge \\ \pi_\alpha(s') = s'' \cdot \text{flush} \wedge (v', v) \in \text{mstate}(s') \end{array} \right) \vee \\ (s = s' \cdot \alpha : v_2 \wedge \pi_\alpha(s') = s'' \cdot \text{load} \wedge (v_1, v_2) \in \text{mstate}(s')) \vee \\ (s = s' \cdot e \wedge e \in \text{inv} \wedge (v_1, v_2) \in \text{mstate}(s')) \end{array} \right)$$

and $\text{mstate}(s) \upharpoonright_i = \{v_i \mid (v_1, v_2) \in \text{mstate}(s)\}$

Basically, the state function `mstate` computes all possibilities of the persisted content v_1 and the buffered content v_2 .

- When crash happens, only the persisted content is preserved and is loaded into the buffered content.
- When storing a value to the cell, the value may be written to the buffered content only or be written to both the persisted and the buffered one.
- When flush happens, the persisted content gets synchronized with the buffered content.
- Any load to the cell gets the buffered content. Moreover, after a load returns, the buffered content is determined, which will eliminate any non-determinism brought up by unflushed stores before previous crashes.

Figure 14 shows the implementation of the FLiT memory cell. To make implementations and proofs more readable, for structure commands like if-statements and loop-statements, we do not unfold them into the encoding using $C_1 + C_2$ and C^* and instead write the code and do the proof in the high-level syntax.

```

1  Mflit:
2  Import M: BCell
3  Import C: Counter
4
5  load() {                               store(v) {
6      v ← M.load();                       C.inc();
7      if (C.get() != 0) {                 M.store(v);
8          M.flush();                       M.flush();
9      }                                     C.dec();
10     return v;                             return;
11 }                                         }
```

Fig. 14. FLiT Memory Cell Implementation

1.1.2 Proof. To prove this example, we need to add more structures to the possibility's definition. We use the possibility with the form of $\rho = (p, s_O)$, where $p \in P_{\text{flit}}$ and is accessed by $\text{lin}(p)$, and s_O is the set of pending invocations. The set P_{flit} is the largest set of atomically linearized traces.

$$p \in P_{\text{flit}} \iff \left(\begin{array}{c} p = \epsilon \vee \\ (p = p' \cdot \alpha : \text{store}(v) \cdot \alpha : \text{ok} \wedge p' \in v_{\text{flit}}) \vee \\ (p = p' \cdot \alpha : \text{load} \cdot \alpha : v \wedge p' \in v_{\text{flit}} \wedge v = \text{fstate}(p')) \end{array} \right)$$

For this definition, we maintain two invariants that: $s \upharpoonright_{\text{FLiT}}$ is linearizable to $p \cdot \langle s_O \rangle$ and $p \cdot \langle s_O \rangle \in v_{\text{flit}}$, where $\langle s_O \rangle$ is any sequence of pending invocations in s_O . We use a ghost variable $B = \text{List } s_O$ to store all buffered pending invocations in order, which implicitly implies that $\forall e \in B. e \in s_O$. As a result, the program state is now a quadruple (Δ, s, ρ, B) . Using this definition of the configuration will not change the soundness of the logic.

There are three states of the system:

- The Flushed state. It means every stores have persisted in the NVM.

$$\text{Flushed} \iff B = \epsilon$$

- The Unflushed state. It means there are buffered stores but the determinism of the buffered content has not been broken by crashes. However, there must be different possibilities in the

persisted content. Buffered stores that are consistent with the current persisted content will directly get linearized.

$$\text{Unflushed} \iff B \neq \epsilon \wedge \left(\forall v_1, v_2 \in \text{mstate}(s \upharpoonright_{M^\epsilon}) \upharpoonright_2. v_1 = v_2 \wedge \right. \\ \left. \exists v_1, v_2 \in \text{mstate}(s \upharpoonright_{M^\epsilon}) \upharpoonright_1. v_1 \neq v_2 \right)$$

- The Unsynced state. It means there are buffered stores and we do not know which one is persisted and loaded into the volatile memory due to some crashes.

$$\text{Unsynced} \iff B \neq \epsilon \wedge \exists v_1, v_2 \in \text{mstate}(s \upharpoonright_{M^\epsilon}) \upharpoonright_2. v_1 \neq v_2$$

The object invariant $I(s, \rho, B)$ is a conjunction of the following conditions. We maintain the invariant at any point of the program.

- The system is always in one of the three states.

$$\text{Flushed} \vee \text{Unflushed} \vee \text{Unsynced}$$

- If the system is in Flushed state, then the memory cell's state is persistent and is the same as the overlay's state.

$$\text{Flushed} \implies \forall (v_1, v_2) \in \text{mstate}(s \upharpoonright_{M^\epsilon}). v_1 = v_2 = \text{fstate}(\rho)$$

- If the system is in Unflushed state, then the memory cell's buffered value is the value of the latest buffered store and any persisted content corresponds to the current overlay's state or a buffered store.

$$\text{Unflushed} \implies \left(\begin{array}{l} (\forall v \in \text{mstate}(s \upharpoonright_{M^\epsilon}) \upharpoonright_2. \text{last}(B \upharpoonright_{\text{store}}) = \text{store}(v)) \wedge \\ (\forall v. (\text{store}(v) \in B \vee \text{fstate}(\rho) = v) \iff v \in \text{mstate}(s \upharpoonright_{M^\epsilon}) \upharpoonright_1) \end{array} \right)$$

- Any Unsynced state is caused by crashes from a Unflushed state. If the system is in Unsynced state, then any content in the memory cell corresponds to the current overlay's state or a buffered store.

$$\text{Unsynced} \implies (\forall v. (\text{store}(v) \in B \vee \text{fstate}(\rho) = v) \iff (v, v) \in \text{mstate}(s \upharpoonright_{M^\epsilon}))$$

- The counter increment by any agent is always non-negative, which implies the counter value to be non-negative, and when the the system is in the Unflushed state, the counter value is non-zero.

$$\forall \alpha. \text{cstate}(\pi_\alpha(s \upharpoonright_{C^\epsilon})) \geq 0 \wedge (\text{Unflushed} \implies \text{cstate}(s \upharpoonright_{C^\epsilon}) \neq 0)$$

Informal explanation of the load proof. Upon invocation of the load to the FLiT memory cell, it invokes the **load to the underlay** M , which can perform three different transitions according to the state.

load-f When the system is in the Flushed state, we directly linearize the load with the underlay's return value, because according to the invariant, we know the underlay's persisted value and buffered value is the overlay's linearized value. Therefore, the linearization of the current load is consistent with the linearized part.

load-uf When the system is in the Unflushed state, we know that v is the underlay's buffered value and according to the invariant, the last buffered store has the argument v . We then append the current load to B , waiting someone to help linearize the load and corresponding store before it to ensure consistency.

load-s When the system is in the Unsynced state, the current load will determine how do buffered operations from previous epochs linearize.

We forget all remaining buffered operation by setting B to empty, because they are no longer visible.

- Or the underlay's load may return the persisted value before all buffered operations, i.e., the same value as the linearized value. Then we forget all buffered operations and linearize the current load after the linearized part.

The post-condition is stabilized to

$$I \wedge \left(\begin{array}{c} (\text{Flushed} \wedge \text{last}(\pi_\alpha(\rho)) = v) \vee \\ ((\text{Unflushed} \wedge (\exists B'. B' \cdot \alpha': \text{store}(v) \cdot \alpha: \text{load} \sqsubseteq B \vee \text{last}(\pi_\alpha(\rho)) = v)) \end{array} \right).$$

The transition load-f and load-us results in the Flushed branch, and load-uf results in the Unflushed branch, where someone may help the current thread linearize its buffered load in B . For simplicity, we use variable v in assertions as the program variable v 's value.

Then, it **gets the current counter value**. In the Flushed branch, the returned counter value n can be any non-negative number. In the Unflushed branch, by invariant I , the counter value must be non-negative. We then stabilize this post-condition into

$$I \wedge \left(\begin{array}{c} (n = 0 \wedge \text{last}(\pi_\alpha(\rho)) = v) \vee \\ (n \neq 0 \wedge (\exists B'. B' \cdot \alpha': \text{store}(v) \cdot \alpha: \text{load} \sqsubseteq B \vee \text{last}(\pi_\alpha(\rho)) = v)) \end{array} \right) \wedge (\text{Flushed} \vee \text{Unflushed})$$

by extracting $\text{Flushed} \vee \text{Unflushed}$ outside, because they can change into each other under the rely. We keep the information that when $n = 0$, the current load is linearized, and otherwise, the current load may remains in B or is linearized.

Into the if branch, we know $n \neq 0$ and therefore, the current load may either be in the buffered list B or already linearized. When it **flushes**, we linearize every thing in the buffered list (which may be empty in the Flushed case), because after the flush, the underlay's persisted value will be consistent with the overlay's value after linearizing operations in B according to the invariant. As a result, after this flush and linearizations, the current load is definitely linearized.

Informal explanation of the store proof. In the load proof, we only use the invariant to get information about the counter, but in the store proof, it is a major challenge to maintain the invariant for the counter. When **increasing the counter** at the beginning of the store, we know from the invariant that the counter increment ($\text{cstate}(\pi_\alpha(s \upharpoonright_{C^i}))$) by the current thread α is non-negative and after the increment, it will be strictly larger than 0 because no one else can change $\text{cstate}(\pi_\alpha(s \upharpoonright_{C^i}))$. This step preserves the invariant, because it does not decrease the counter value.

Then it **performs a store**, which can have three different transitions under four different cases.

store-ff If we are in the Flushed state and we are storing a value that is exactly the same as the current linearized value. Then we can directly linearize this store because the underlay is the same no matter this store is persisted or not. Another choice is to put it in the buffered list and let someone else linearize it, but this makes defining the Unflushed states difficult. This step preserves the invariant because we are still in the Flushed state.

store-fu If we are in the Flushed state and we are storing a value that is different from the current linearized value. Then we need to put it in the buffered list, because it may not be persisted and is unsafe to be linearized. This step preserves the invariant, because although we changed into the Unflushed state, we have ensured the local increment to the counter is positive, which means the counter value is positive.

store-uu (unflushed) If we are in the Unflushed state, we append the store to the buffered list B and it is still in the Unflushed state. This step preserves the invariant for the same reason as the store-fu transition.

```

{invokeα(store(v)) ◦ I}
1: store(v){
  {I ∧ α:store(v) ∈ sO ∧ 0 ≤ cstate(πα(s↑Ci))}
2: C.inc();
  {I ∧ α:store(v) ∈ sO ∧ 0 < cstate(πα(s↑Ci)) ∧ (Flushed ∨ Unflushed ∨ Unsynced)}
3: M.store(v); // store-ff/store-fu/store-uu
  {
    I ∧ 0 < cstate(πα(s↑Ci)) ∧ s' = s · α:M.store(v) · α:ok ∧
    {
      (Flushed(s, B) ∧ fstate(ρ) = v ∧ B' = ε ∧ lin(ρ') = lin(ρ) · α:store(v) · α:ok) ∨
      (Flushed(s, B) ∧ fstate(ρ) ≠ v ∧ B' = α:store(v) · ε ∧ ρ' = ρ) ∨
      ((Unflushed(s, B) ∨ Unsynced(s, B)) ∧ B' = B · α:store(v) ∧ ρ' = ρ)
    }
    {
      I ∧ 0 < cstate(πα(s↑Ci)) ∧
      ((Unflushed ∧ last(πα(ρ)) = ok) ∨
      (Flushed ∧ last(πα(ρ)) = ok))
    }
  }
4: M.flush(); // flush
  {
    I ∧ 0 < cstate(πα(s↑Ci)) ≤ cstate(s↑Ci)) ∧
    {
      (Flushed(s, B) ∧ last(πα(ρ)) = ok ∧ (ρ', B') = (ρ, B)) ∨
      {
        (Unflushed(s, B) ∧ lin(ρ') = merge(lin(ρ), B) ∧
        B' = ε ∧ ((α:store(v) ∈ B ∨ last(πα(ρ)) = ok))
      }
    }
    {
      I ∧ last(πα(ρ)) = ok ∧
      ((Flushed ∧ 0 < cstate(πα(s↑Ci)) ≤ cstate(s↑Ci)) ∨
      (Unflushed ∧ 0 < cstate(πα(s↑Ci)) < cstate(s↑Ci)))
    }
5: C.dec();
  {
    I ∧ last(πα(ρ)) = ok ∧
    ((Flushed ∧ 0 ≤ cstate(πα(s↑Ci)) ≤ cstate(s↑Ci)) ∨
    (Unflushed ∧ 0 ≤ cstate(πα(s↑Ci)) < cstate(s↑Ci)))
  }
6: ret ok
7: }
{returnedα(store(v)) ◦ I} {⊤}

```

Fig. 16. Proof of the FLiT Memory Cell: store

store-uu (unsynced) If we are in the Unsynced state, we append the store to the buffered list B and step into the Unflushed state. This transition is valid because an underlay store synchronizes all possible buffered values to be one value after the crash breaks such synchronization. This step preserves the invariant for the same reason as the store-fu transition.

We stabilize the post-condition into

$$I \wedge 0 < \text{cstate}(\pi_\alpha(s \uparrow_{C^i})) \wedge \left(\begin{array}{c} (\text{Flushed} \wedge \text{last}(\pi_\alpha(\rho)) = \text{ok}) \vee \\ (\text{Unflushed} \wedge ((\alpha:\text{store}(v) \in B \vee \text{last}(\pi_\alpha(\rho)) = \text{ok}))) \end{array} \right)$$

The transition store-ff results into the Flushed branch and other two transitions results into the Unflushed branch. Other threads may change Flushed into Unflushed but the current store will remain linearized or change Unflushed into Flushed and guarantee the current store will be linearized since it is in B .

Then it **flushes**. By definition, the counter value ($\text{cstate}(s \uparrow_{C^i})$) is no less than the local increment ($\text{cstate}(\pi_\alpha(s \uparrow_{C^i}))$). In the Flushed case, the flush operation has no effect. In the Unflushed case, the flush operation will linearize all buffered operations in B , which includes the current store if it has not been linearized. Therefore, it ensures the current store is linearized and results in the

Flushed in the stabilized post-condition:

$$I \wedge \text{last}(\pi_\alpha(\rho)) = \text{ok} \wedge \left((\text{Flushed} \wedge 0 < \text{cstate}(\pi_\alpha(s \upharpoonright_{C^i})) \leq \text{cstate}(s \upharpoonright_{C^i})) \vee \right. \\ \left. (\text{Unflushed} \wedge 0 < \text{cstate}(\pi_\alpha(s \upharpoonright_{C^i})) < \text{cstate}(s \upharpoonright_{C^i})) \right)$$

This step preserves the invariant because it changes into the Flushed state. However, other threads may change it into the Unflushed state. Luckily, when they make the store-fu or store-uu transition, they will guarantee that their local counter increment is non-zero, which means the counter value is strictly larger than the local increment of the current thread. This fact is important when proving the next step.

Then it **decreases the counter**. This step does not change the system state, but we need to verify that it preserves the invariant. In either branch, the local increment ($\text{cstate}(\pi_\alpha(s \upharpoonright_{C^i}))$) now is non-negative. If the system is still Flushed, the invariant is preserved. If the system is Unflushed, since we have maintained that $\text{cstate}(\pi_\alpha(s \upharpoonright_{C^i})) < \text{cstate}(s \upharpoonright_{C^i})$, we know the counter value is still non-zero and the invariant is preserved.

1.1.3 Summary & More Proof Details. To summarize, we define the invariant I of the object as the following, which is satisfied at any point of the program and is the crash-invariant as well.

$$I(s, \rho, B) \iff \left(\begin{array}{l} (\text{Flushed} \vee \text{Unflushed} \vee \text{Unsynced}) \wedge \\ (\text{Flushed} \Rightarrow \forall (v_1, v_2) \in \text{mstate}(s \upharpoonright_{M^i}). v_1 = v_2 = \text{fstate}(\rho)) \wedge \\ (\text{Unflushed} \Rightarrow \left(\begin{array}{l} (\forall v \in \text{mstate}(s \upharpoonright_{M^i}) \upharpoonright_2. \text{last}(B \upharpoonright_{\text{store}}) = \text{store}(v)) \wedge \\ (\forall v. (\text{store}(v) \in B \vee \text{fstate}(\rho) = v) \Leftrightarrow v \in \text{mstate}(s \upharpoonright_{M^i}) \upharpoonright_1) \end{array} \right)) \wedge \\ (\text{Unsynced} \Rightarrow (\forall v. (\text{store}(v) \in B \vee \text{fstate}(\rho) = v) \Leftrightarrow (v, v) \in \text{mstate}(s \upharpoonright_{M^i}))) \wedge \\ \forall \alpha. \text{cstate}(\pi_\alpha(s \upharpoonright_{C^i})) \geq 0 \wedge (\text{Unflushed} \Rightarrow \text{cstate}(s \upharpoonright_{C^i}) \neq 0) \end{array} \right)$$

And there are three states of the object with transitions illustrated in Fig. 17.

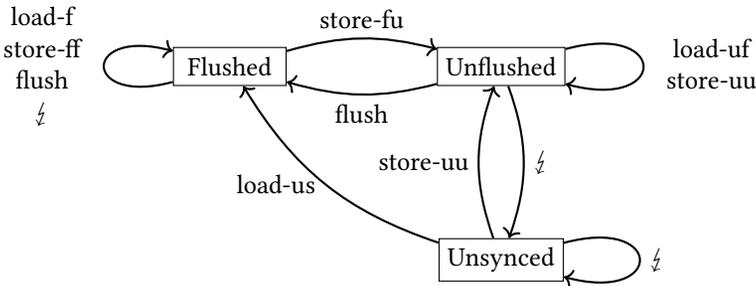


Fig. 17. STS for FLiT Memory Cell

Notice that this invariant I is idempotent under the crash-into relation, i.e., $I \Rightarrow_{\zeta} I$.

- Firstly, as the state transition system in Fig. 17 and the invariant definition indicates, the crash will change the system into the Flushed or Unsynced state, which is still in the invariant.
- Secondly, a crash does not invalidate any conjuncts in the invariant. It is worth mentioning that when a crash changes Unflushed in Unsynced, the second branch of the Unflushed branch, which is about the persisted value of M , is necessary for establishing the invariant in the Unsynced branch.

This object does not need a recovery program, so we set the body of r to be `ret ok`, which does nothing and returns immediately. And it is obvious that the following quadruple holds.

$$\text{ID}, \top \vDash_{\alpha} \{\text{invoke}_{\alpha}(r) \circ I\} \text{ret ok} \{\text{returned}_{\alpha}(r) \circ I\} \{I\}$$

The pre-/post-conditions, I , of this recovery program obviously connect to the crash invariant and the object invariant, which are also I .

We have proved $\mathcal{R}[\alpha], \mathcal{G}[\alpha], I[\alpha] \vDash_{\alpha} M_F[A]$ in the previous section. And it is easy to check that we have established all conditions (except the stability) in OBJECT IMPL for the judgement

$$P_{\dagger}(\text{FLIT} \cup R_{\phi}) \vdash M_{\text{FLIT}} : (v'_{\text{BCell}} \boxtimes v'_{\text{Counter}}, v_{\text{BCell}} \otimes v_{\text{Counter}}) \rightarrow \langle v'_{\text{FLIT}}, v_{\text{FLIT}} \rangle$$

We then define the guarantee relations of each transitions and they form the rely relation. It is easy to check that all stabilized assertions are stable w.r.t. the rely relation. And this finishes the proof of the FLIT memory cell.

$$\begin{aligned} (s, \rho, B) \mathcal{G}_{\text{load-f}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \exists v. \text{Flushed}(s, B) \wedge (v, v) \in \text{mstate}(s \upharpoonright_{M^i}) \wedge v = \text{fstate}(\rho) \wedge \\ \text{lin}(\rho') = \text{lin}(\rho) \cdot \alpha : \text{load} \cdot \alpha : v \wedge B' = \epsilon \wedge \\ s' = s \cdot \alpha : M.\text{load} \cdot \alpha : v \end{array} \right) \\ (s, \rho, B) \mathcal{G}_{\text{load-uf}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \exists v. \text{Unflushed}(s, B) \wedge \text{last}(B \upharpoonright_{\text{store}}) = \text{store}(v) \wedge \\ B' = B \cdot \alpha : \text{load} \wedge \rho' = \rho \wedge \\ s' = s \cdot \alpha : M.\text{load} \cdot \alpha : v \end{array} \right) \\ (s, \rho, B) \mathcal{G}_{\text{load-us}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \exists v, B_1, B_2. \text{Unsynced}(s, B) \wedge (v, v) \in \text{mstate}(s \upharpoonright_{M^i}) \wedge \\ \left((B = B_1 \cdot B_2 \wedge \text{last}(B_1) = \text{store}(v)) \right) \wedge \\ \vee (\text{fstate}(\rho) = v \wedge B_1 = \epsilon) \\ \text{lin}(\rho') = \text{merge}(\text{lin}(\rho), B_1) \cdot \alpha : \text{load} \cdot \alpha : v \wedge B' = \epsilon \wedge \\ s' = s \cdot \alpha : M.\text{load} \cdot \alpha : v \end{array} \right) \\ (s, \rho, B) \mathcal{G}_{\text{store-ff}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \exists v. \text{Flushed}(s, B) \wedge 0 < \text{cstate}(\pi_{\alpha}(s \upharpoonright_{C^i})) \wedge \text{fstate}(\rho) = v \wedge \\ \text{lin}(\rho') = \text{lin}(\rho) \cdot \alpha : \text{store}(v) \cdot \alpha : \text{ok} \wedge B' = \epsilon \\ s' = s \cdot \alpha : M.\text{store}(v) \cdot \alpha : \text{ok} \end{array} \right) \\ (s, \rho, B) \mathcal{G}_{\text{store-fu}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \exists v. \text{Flushed}(s, B) \wedge 0 < \text{cstate}(\pi_{\alpha}(s \upharpoonright_{C^i})) \wedge \text{fstate}(\rho) \neq v \wedge \\ B' = \alpha : \text{store}(v) \cdot \epsilon \wedge \rho' = \rho \\ s' = s \cdot \alpha : M.\text{store}(v) \cdot \alpha : \text{ok} \end{array} \right) \\ (s, \rho, B) \mathcal{G}_{\text{store-uu}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \exists v. (\text{Unflushed}(s, B) \vee \text{Unsynced}(s, B)) \wedge \\ 0 < \text{cstate}(\pi_{\alpha}(s \upharpoonright_{C^i})) \wedge B' = B \cdot \alpha : \text{store}(v) \wedge \\ \rho' = \rho \wedge s' = s \cdot \alpha : M.\text{store}(v) \cdot \alpha : \text{ok} \end{array} \right) \end{aligned}$$

$$\begin{aligned}
(s, \rho, B) \mathcal{G}_{\text{flush}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \text{lin}(\rho') = \text{merge}(\text{lin}(\rho), B) \wedge B' = \epsilon \wedge \\ s' = s \cdot \alpha:\text{flush} \cdot \alpha:\text{ok} \end{array} \right) \\
(s, \rho, B) \mathcal{G}_{\text{inc}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \rho' = \rho \wedge B' = B \wedge s' = s \cdot \alpha:\text{inc} \cdot \alpha:\text{ok} \wedge \\ \text{cstate}(\pi_\alpha(s' \upharpoonright_{C^i})) = \text{cstate}(\pi_\alpha(s \upharpoonright_{C^i})) + 1 \end{array} \right) \\
(s, \rho, B) \mathcal{G}_{\text{dec}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \rho' = \rho \wedge B' = B \wedge s' = s \cdot \alpha:\text{dec} \cdot \alpha:\text{ok} \wedge \\ \text{cstate}(\pi_\alpha(s' \upharpoonright_{C^i})) = \text{cstate}(\pi_\alpha(s \upharpoonright_{C^i})) - 1 \end{array} \right) \\
(s, \rho, B) \mathcal{G}_{\text{id}}[\alpha](s', \rho', B') &\iff \left(\begin{array}{l} \rho' = \rho \wedge B' = B \wedge \\ \text{mstate}(\pi_\alpha(s' \upharpoonright_{M^i})) = \text{mstate}(\pi_\alpha(s \upharpoonright_{M^i})) \wedge \\ \text{cstate}(\pi_\alpha(s' \upharpoonright_{C^i})) = \text{cstate}(\pi_\alpha(s \upharpoonright_{C^i})) \end{array} \right) \\
\mathcal{R}[\alpha] \triangleq \bigcup_{\alpha' \in \Upsilon, \alpha' \neq \alpha} &\left(\begin{array}{l} \mathcal{G}_{\text{load-f}}[\alpha'] \cup \mathcal{G}_{\text{load-uf}}[\alpha'] \cup \mathcal{G}_{\text{load-us}}[\alpha'] \cup \mathcal{G}_{\text{store-ff}}[\alpha'] \cup \mathcal{G}_{\text{store-fu}}[\alpha'] \cup \\ \mathcal{G}_{\text{store-uu}}[\alpha'] \cup \mathcal{G}_{\text{flush}}[\alpha'] \cup \mathcal{G}_{\text{inc}}[\alpha'] \cup \mathcal{G}_{\text{dec}}[\alpha'] \cup \mathcal{G}_{\text{id}}[\alpha'] \cup \\ \text{invoke}_{\alpha'}(-) \cup \text{return}_{\alpha'}(-) \end{array} \right)
\end{aligned}$$

1.2 The Interval-Sequential Write-Snapshot Object

With the FLiT memory cell's correctness established in appendix I.1 and the FLiT correctness theorem (proposition 1.1), we can prove any object implement to be durably linearizable as long as we can prove it satisfying linearizability defined in [31]. For example, the interval-sequential write-snapshot object in [7] can be implemented as a durably linearizable object using the FLiT cell's read and write instead of the usual atomic memory cell. In this section, we prove the interval-linearizability of the one-shot write-snapshot object in [7] using the program logic in [31].

1.2.1 Specification & Implementation. The write-snapshot object has the signature below.

$$\text{Snapshot} := \{\text{write_snapshot} : \text{Val} \rightarrow 2^{\text{Val}}\}$$

The operation `write_snapshot` writes the current value to the memory and returns a set of values that have been written to the object before.

Its interval-sequential specification $\nu_{\text{write_snapshot}}$ is the largest set of traces satisfying the following

$$s \in \nu_{\text{write_snapshot}} \iff \left(\begin{array}{l} s = \epsilon \\ \vee (s = s' \cdot \alpha:V \wedge V = \text{snpstate}(s') \wedge s' \in \nu_{\text{write_snapshot}}) \\ \vee (s = s' \cdot \alpha:\text{write_snapshot}(v) \wedge \alpha \notin s' \wedge s' \in \nu_{\text{write_snapshot}}) \end{array} \right)$$

where the state of the write-snapshot object $\text{snpstate}(s)$ is defined as the set of values given by all invocations before a point.

$$\text{snpstate}(s) := \begin{cases} \{\} & \text{if } s = \epsilon \\ \text{snpstate}(s') \cup \{v\} & \text{if } s = s' \cdot \alpha:\text{write_snapshot}(v) \\ \text{snpstate}(s') & \text{otherwise, } s = s' \cdot \alpha : _ \end{cases}$$

The specification clarifies two key points of the object:

- Firstly, the object is not atomic-sequential, which means the linearized specification does not guarantee the response happens immediately after the corresponding invocation. The response will take into account all invocations linearized before itself.

- Secondly, the object is one-shot. Each agent (thread) can only invoke $v_{\text{write_snapshot}}$ once. We use the client specification $\mu_{\text{write_snapshot}}$ to pose this one-shot requirement on clients.

$$s \in \mu_{\text{write_snapshot}} \iff \left(\begin{array}{l} \forall s', \alpha, v, s' \cdot \alpha : \text{write_snapshot}(v) \sqsubseteq s \Rightarrow \\ (\forall v, \alpha : \text{write_snapshot}(v') \notin s' \wedge \alpha \in S) \end{array} \right)$$

As [7], we consider the one-shot write-snapshot algorithm for solving problems with only finite participating agents. Therefore, we require the client to take a finite subset S of all agents Υ , and use the client specification to require all invocations be made in these threads.

The implementation of the write-snapshot object is shown in figure 18. With a total number of $|S|$ threads where the write-snapshot object will be used, each thread is assigned a FLiT memory cell (with initial value \perp) to store the value it writes to the object. Since this is a one-shot object, we know that each memory cell is written at most once. After written to the cell, the algorithm repeatedly takes a snapshot of all values written to the object. It returns only when two consecutive snapshots converges, which guarantees the snapshot to be correct, meaning the snapshot did occur at a certain moment. This is non-trivial. For example, when taking a snapshot, if two writes to a visited cell and un-visited cell occurs, only the second one will be captured in the snapshot, which makes the snapshot invalid because in any possible snapshot, the second value should be recorded if the first value is not recorded.

```

1   $M_{\text{write\_snapshot}}$  :
2  Import M:  $\otimes_{i \in |S|}$  FLiTi
3
4  write_snapshot(int v) {
5    M[ $\alpha$ ].write(v);
6    old  $\leftarrow$  { $\perp$ }; new  $\leftarrow$   $\emptyset$ ; i  $\leftarrow$  1;
7    while (i  $\leq$  |S|) {
8      v  $\leftarrow$  M[ $\alpha_i$ ].read();
9      new  $\leftarrow$  new  $\cup$  {v};
10     i  $\leftarrow$  i+1
11   }
12   while (new  $\neq$  old) {
13     old  $\leftarrow$  new; new  $\leftarrow$   $\emptyset$ ; i  $\leftarrow$  1;
14     while (i  $\leq$  |S|) {
15       v  $\leftarrow$  M[ $\alpha_i$ ].read();
16       new  $\leftarrow$  new  $\cup$  {v};
17       i  $\leftarrow$  i+1
18     }
19   };
20   return new \ { $\perp$ }
21 }

```

Fig. 18. One-Shot Write-Snapshot Implementation

1.2.2 Proof. Like what we did in the FLiT example, we use the possibility with the form of $\rho = (p, s_O)$, where p is the linearized trace with $p \in v_{\text{write_snapshot}}$ (and is accessed through $\text{lin}(\rho)$) and s_O are pending invocations. Notice that there can exist pending invocations in p as well, since we are proving an interval-sequential object which allows an interval without the response, but these invocations will have impact on linearized responses in the future.

We do not use any other ghost variables in this proof, and the program configuration is a triple (Δ, s, ρ) . We need to maintain two invariants that: $s \upharpoonright_{\text{Snapshot}}$ is linearizable to $p \cdot \langle s_O \rangle$ and $p \cdot \langle s_O \rangle \in v_{\text{write_snapshot}}$.

For this concurrent object, we maintain the following object invariant at any point of the program execution. It simply ensures the the order of writes in s is the same as the order of `write_snapshot` in the linearized trace.

$$I(\Delta, s, \rho) \iff s \upharpoonright_{\text{write}} \sim \text{lin}(\rho) \upharpoonright_{\text{write_snapshot}}$$

$$\text{where } s_1 \sim s_2 \iff \left(\begin{array}{c} (s_1 = s_2 = \epsilon) \vee \\ \exists \alpha, v, s'_1, s'_2. \left(\begin{array}{c} s'_1 \sim s'_2 \wedge s_1 = s'_1 \cdot \alpha : \text{write}(v) \\ \wedge s_2 = s'_2 \cdot \alpha : \text{write_snapshot}(v) \end{array} \right) \end{array} \right)$$

For the three while loop, we introduce the following two loop invariant.

$$I_o[i, \text{old}, p](\Delta, s, \rho) \iff p \sqsubseteq \text{lin}(\rho) \wedge \text{old} \setminus \{\perp\} \subseteq \bigcup_{j=1}^{i-1} \text{snpstate}(\pi_{\alpha_j}(p)) \wedge i \leq |S| + 1$$

$$I_n[i, \text{new}, p](\Delta, s, \rho) \iff \left(\begin{array}{c} \bigcup_{j=1}^{i-1} \text{snpstate}(\pi_{\alpha_j}(p)) \subseteq \text{new} \setminus \{\perp\} \wedge i \leq |S| + 1 \\ \wedge (\forall \alpha : \text{write_snapshot}(v) \in p \Rightarrow \alpha : M[\alpha].\text{write}(v) \in s) \end{array} \right)$$

- The loop invariant I_o guarantees there is always a prefix p of the linearized possibility such that each write in the old snapshot is captured by p . The invariant I_o asserts that the old snapshot is a lower bound of a linearized trace's snapshot state at a certain moment.
- The loop invariant I_n guarantees that any write in the prefix p extracted from I_o is always in the new snapshot. The invariant I_n asserts that the new snapshot is an upper bound of the linearized trace p 's snapshot state.
- Both invariant takes a parameter i , which is the loop variable. We will only consider the first i threads operations in the trace and maintain the invariant. When the loop terminates, i becomes $|S| + 1$, and the invariant will be true under the consideration of all threads, which means it is true for the complete trace.

When the outer loop terminates, we know $\text{old} \setminus \{\perp\} = \text{snpstate}(p) = \text{new} \setminus \{\perp\}$, and we can linearize the response at the end of p since $\text{new} \setminus \{\perp\}$ is valid snapshot at that moment.

```

{invokeα(write_snapshot) ◦ I}
1: write_snapshot(v){
  {I ∧ α:write_snapshot(v) ∈ sO}
2: M[α].write(v); // write
  {I ∧ last(πα(lin(ρ))) = α:write_snapshot(v)}
3: old ← {⊥}; new ← ∅; i ← 1;
  {I ∧ last(πα(lin(ρ))) = α:write_snapshot(v) ∧ Io[|S| + 1, old, ε] ∧ In[i, new, ε] ∧ ∃q. Io[i, new, q]}
  {I ∧ last(πα(lin(ρ))) = α:write_snapshot(v) ∧ ∃p. Io[|S| + 1, old, p] ∧ In[i, new, p] ∧ ∃q. Io[i, new, q]}
4: while(i ≤ |S|){
5:   v ← M[αi].read();
6:   new ← new ∪ {v}; i ← i + 1
7: }
  {
    I ∧ last(πα(lin(ρ))) = α:write_snapshot(v) ∧
    {∃p. Io[|S| + 1, old, p] ∧ In[|S| + 1, new, p] ∧ ∃q. Io[|S| + 1, new, q]}
8: while(new ≠ old){
  {I ∧ last(πα(lin(ρ))) = α:write_snapshot(v) ∧ ∃q. Io[|S| + 1, new, q]}
9:   old ← new; new ← ∅; i ← 1;
  {I ∧ last(πα(lin(ρ))) = α:write_snapshot(v) ∧ ∃p. Io[|S| + 1, old, p] ∧ In[i, new, p] ∧ ∃q. Io[i, new, q]}
10:  while(i ≤ |S|){
11:    v ← M[αi].read();
12:    new ← new ∪ {v}; i ← i + 1
13:  }
  {
    I ∧ last(πα(lin(ρ))) = α:write_snapshot(v) ∧
    {∃p. Io[|S| + 1, old, p] ∧ In[|S| + 1, new, p] ∧ ∃q. Io[|S| + 1, new, q]}
14: }
  {I ∧ last(πα(lin(ρ))) = α:write_snapshot(v) ∧ ∃p. Io[|S| + 1, new, p] ∧ In[|S| + 1, new, p]}
  {I ∧ last(πα(lin(ρ))) = α:write_snapshot(v) ∧ ∃p ⊆ lin(ρ, L). snpstate(p) = new \ {⊥}}
  {I ∧ last(πα(ρ)) = α:new\{⊥}} // snapshot
15: ret new \ {⊥}
16: }
{returnedα(write_snapshot) ◦ I}

```

Fig. 19. Proof of Write-Snapshot

Figure 19 shows the proof outline for the write-snapshot object. At line 2, when writing to the FLiT cell, we linearize the corresponding write_snapshot invocation to the linearized trace so that the object invariant can be maintained. This operation satisfies the following guarantee $\mathcal{G}_{\text{write}}$.

$$(\Delta, s, \rho) \mathcal{G}_{\text{write}}[\alpha](\Delta', s', \rho') \iff \left(\begin{array}{l} s' = s \cdot \alpha : M[\alpha]\text{write}(v) \cdot \alpha : \text{ok} \\ \wedge s'_O = s_O \setminus \{\alpha:\text{write_snapshot}(v)\} \\ \wedge \text{lin}(\rho') = \text{lin}(\rho) \cdot \alpha : \text{write_snapshot}(v) \end{array} \right)$$

Then, we initial the old snapshot old, new snapshot new, and the loop variable i . Since old contains only the initial value, it is a lower bound of any trace's snapshot state, and therefore we can take $p = \epsilon$ and $I[|S| + 1, \text{old}, \epsilon]$ is true. Moreover, since ϵ produce empty snapshot state, any set is an upper bound of it and thus $I_n[i, \text{new}, \epsilon]$ is true for any i .

The loop from line 4 to line 6 is the process of taking the snapshot. We may split the proof for this loop into two parts,

- one for constructing I_o , the lower bound invariant for the next snapshot loop;
- one for constructing I_n , the upper bound invariant for the current snapshot loop.

Their proofs are independent, and since there are no control flow can break out of the loop, we may present their proofs separately for a clearer demonstration.

Construct Lower Bound. In figure 20, we use $I \wedge \exists p. I_o[i, \text{new}, p]$ as the loop invariant. In each iteration, the read operation does not change the state of the memory cell and the write-snapshot object and there are two cases for the read value.

- If the value is \perp , meaning thread α_i has not written to the cell and new is still the lower bound, i.e., $I_o[i + 1, \text{new} \cup \{\perp\}, p]$ is true.
- If the value is not \perp , then α_i has written to its cell, and by the object invariant I , its `write_snapshot(v)` invocation has been linearized. Therefore, there exists $q \sqsubseteq \text{lin}(\rho)$ which includes this invocation and we can take the longest one of p and q . Since this object is one-shot, values written to any cell cannot be overwrite when we consider more future events by extending p to $\max\{p, q\}$, and any value in new still appears in $\bigcup_{j=1}^{i-1} \text{snpstate}(\pi_{\alpha_j}(\max\{p, q\}))$. As a result, $\text{new} \cup \{v\}$ is a lower bound considering only the first $i - 1$ threads, i.e., $I_o[i + 1, \text{new} \cup \{v\}, \max\{p, q\}]$ is true.

Then we can safely include v in new and increase the loop counter while maintaining the invariant.

```

{I ∧ ∃p. I_o[i, new, p]}
1: while(i ≤ |S|){
  {I ∧ ∃p. I_o[i, new, p] ∧ i ≤ |S|}
2: v ← M[α_i].read();
  {
    I ∧ ∃p. I_o[i, new, p] ∧ i ≤ |S|
    ∧ (
      (v = ⊥ ∧ I_o[i + 1, new ∪ {⊥}, p])
      ∨ (
        v ≠ ⊥ ∧ α_i:M[α_i].write(v) ∈ s
        ∧ α_i:write_snapshot(v) ∈ lin(ρ)
        ∧ ∃q ⊆ lin(ρ). α_i:write_snapshot(v) ∈ q
        ∧ I_o[i + 1, new ∪ {v}, max{p, q}]
      )
    )
  }
3: new ← new ∪ {v}; i ← i + 1
  {I ∧ ∃p. I_o[i, new, p]}
4: }
  {I ∧ ∃p. I_o[|S| + 1, new, p]}

```

Fig. 20. Lower Bound Construction Proof

Construct Upper Bound. In figure 21, we use $I \wedge \exists p. I_o[|S| + 1, \text{old}, p] \wedge I_n[i, \text{new}, p]$ as the loop invariant. According to the read value, there will be four different cases.

- If the read value is \perp , then thread α_i has not written to its cell in $\text{lin}(\rho)$. Obviously, its write will no appear in the prefix p and $\text{new} \cup \{\perp\}$ is an upper bound, i.e., $I_n[i + 1, \text{new} \cup \{\perp\}, p]$ is true.
- If the read value is not \perp and $\alpha_i:\text{write_snapshot}(v)$ already appears in the prefix p , then it is safe to include v in the upper bound, i.e., $I_n[i + 1, \text{new} \cup \{v\}, p]$ is true.
- If the read value is not \perp and $\alpha_i:\text{write_snapshot}(w)$ does not appear in the prefix p for any w , then we can safely add v to the upper bound for the same reason as the first case, and $I_n[i + 1, \text{new} \cup \{v\}, p]$ is true.
- If the read value is not \perp and $\alpha_i:\text{write_snapshot}(w)$ already appears in the prefix p for a different w from v , then adding v to the upper bound new may produce a snapshot inconsistent with the trace. However, since we read a different v from the memory cell, $\alpha_i:\text{write}(v)$

must appear in s , and by I , we know $\alpha_i:\text{write_snapshot}(v)$ must appear in $\text{lin}(\rho)$. With two different write_snapshot invocations for the same thread α_i in $\text{lin}(\rho)$, the client specification $\mu_{\text{write_snapshot}}$ is violated, which leads to a contradiction with $\mu_{\text{write_snapshot}}$. As a result, under the restriction of $\mu_{\text{write_snapshot}}$, this case will not happen.

Then we can safely include v in new and increase the loop counter while maintaining the invariant.

```

{I ∧ ∃p. I_o[|S| + 1, old, p] ∧ I_n[i, new, p]}
1: while(i ≤ |S|){
  {I ∧ ∃p. I_o[|S| + 1, old, p] ∧ I_n[i, new, p] ∧ i ≤ |S|}
2: v ← M[α_i].read();
   {
     I ∧ ∃p. I_o[|S| + 1, old, p] ∧ I_n[i, new, p] ∧ i ≤ |S|
     (v = ⊥ ∧ I_n[i + 1, new ∪ {⊥}, p])
     ∨ (v ≠ ⊥ ∧ α_i:write_snapshot(v) ∈ p ∧ I_n[i + 1, new ∪ {v}, p])
     ∨ (v ≠ ⊥ ∧ ∀w. α_i:write_snapshot(w) ∉ p ∧ I_n[i + 1, new ∪ {v}, p])
     ∨ (v ≠ ⊥ ∧ ∃w ≠ v. α_i:write_snapshot(w) ∈ p ∧ α_i:M[α_i].write(v) ∈ s)
     ∧ α_i:write_snapshot(v) ∈ lin(ρ) ∧ α_i:write_snapshot(w) ∈ lin(ρ)
   }
3: new ← new ∪ {v}; i ← i + 1
   {I ∧ ∃p. I_o[|S| + 1, old, p] ∧ I_n[i, new, p]}
4: }
{I ∧ ∃p. I_o[|S| + 1, old, p] ∧ I_n[|S| + 1, new, p]}

```

Fig. 21. Upper Bound Construction Proof

Then, we reach the loop at line 8 in figure 19. Here, we use the loop invariant

$$I \wedge \text{last}(\pi_\alpha(\text{lin}(\rho))) = \alpha:\text{write_snapshot}(v) \wedge \\ \exists p. I_o[|S| + 1, \text{old}, p] \wedge I_n[|S| + 1, \text{new}, p] \wedge \exists q. I_o[|S| + 1, \text{new}, q]$$

which asserts that old and new are lower and upper bounds of the snapshot state for some prefix p of the linearized trace, and new is the lower bound of another linearized prefix q .

- Inside the loop, we only keep the second branch and by assigning new to old , the new old is still the lower bound. Then, a loop identical to the one at line 4 takes a new snapshot into new and we reuse the previous proof for it. The loop establishes the new as the new upper bound for the new prefix p and we re-establish the loop invariant.
- If the loop terminates, we only keep the first branch. With the snapshot state bounded by new on both side, we can derive $\text{snpstate}(p) = \text{new} \setminus \{\perp\}$, which means $\text{new} \setminus \{\perp\}$ is the correct snapshot at the time where p is the complete linearized trace. We can then linearize the response to the current operation after p and $\text{last}(\pi_\alpha(\rho)) = \alpha:\text{new} \setminus \{\perp\}$ is true after the linearization, which finishes the proof. This linearization step satisfies the following guarantee.

$$(\Delta, s, \rho) \mathcal{G}_{\text{snapshot}}[\alpha](\Delta', s', \rho') \iff \left(\begin{array}{l} s' = s \wedge \exists p_1, p_2, V. \text{lin}(\rho) = p_1 \cdot p_2 \\ \wedge \text{lin}(\rho') = p_1 \cdot \alpha:V \cdot p_2 \wedge V = \text{snpstate}(p_1) \end{array} \right)$$

The rely rely is defined as

$$\mathcal{R}[\alpha] \triangleq \bigcup_{\alpha' \in \Upsilon, \alpha' \neq \alpha} \mathcal{G}_{\text{write}}[\alpha'] \cup \mathcal{G}_{\text{snapshot}}[\alpha'] \cup \mathcal{G}_{\text{id}}[\alpha'] \cup \text{invoke}_{\alpha'}(-) \cup \text{return}_{\alpha'}(-)$$

and one may easily check that all assertions in the proof is stable w.r.t. it. Now, we have proved that the write-snapshot object is linearizable w.r.t. its specification $\nu_{\text{write_snapshot}}$ and by the FLiT correctness theorem, this object is also durably linearizable.

I.3 Swap Operation in File System through Write-Ahead Logs

In this four-layered example, we demonstrate that our framework can handle sophisticated file system patterns compositionally. In the upper layer, we present a file system capable of swapping files atomically between directories. The file system depends on a write-ahead log objects and an array of replicated disk cells, both of which are implemented in the lower layer. The file system is presented in I.3.1, while the replicated disks and write-ahead log is described in I.3.3 and I.3.5.

I.3.1 File System with Read, Write, and Swap. We first present and verify a file system that supports file read and write operation, as well as file swapping between directories. The file system exposes a two level structure. At the first level lies a set of folders, each occupies a single disk block as their inode. For simplicity, the API uses block ids instead of strings to uniquely identify folders. Each folder contains a set of files, identified by their file id (unique within each folder). A swap operation will swap the pointer in respective folders' inodes, which can be considered as a symmetric move operation seen in actual file systems. For simplicity, we restrict each file to contain a single block for file content, and assumes all files and directories are pre-allocated on the disk. Allowing for transactions involving several blocks is straight-forward given that we show our techniques work in this simplified setting.

While file read and write operations are mostly straightforward, the swap operation requires special treatment for its atomicity. As swap operations need to update two different folders (and thus two different disk blocks), the possibility that a crash happens in between can never be ruled out. To ensure persistent linearizability, we record the operations in write-ahead logs so that the recovery routine can finish incomplete operations. Figure 22 showcases the pseudocode for this file system.

The underlying disk object is modeled as a map from `block_id` to `block`. `block_id` can be considered as an integer type whereas `block` can be considered as a constant-sized byte array. In the case a block actually contains folder inode information, we byte-cast them into the correct type `folder_inode`.

$$\begin{aligned}
 \text{Disk} &:= \{\text{write} : \text{block_id} \times \text{block} \rightarrow 1, \text{read} : \text{block_id} \rightarrow \text{block}\} \\
 \text{eval}_{\text{Disk}^\xi} &: P_{\dagger\text{Disk}^\xi} \rightarrow \{\perp\} + (\text{file_id} \xrightarrow{\text{fin}} \text{block}) \\
 \text{eval}_{\text{Disk}^\xi}(s) &:= \begin{cases} \bigcup_{b \in \text{block_id}} \{b \leftarrow 0\} & \epsilon \\ M & s = s' \cdot \xi \wedge \text{eval}_{\text{Disk}^\xi}(s') = M \\ M & s = s' \cdot m \cdot \xi \wedge \text{eval}_{\text{Disk}^\xi}(s' \cdot \xi) = M \wedge \lambda_{\text{Disk}}(m) = O \\ M[b \mapsto v] & s = s' \cdot \alpha:\text{write}(b, v) \cdot \alpha:\text{ok} \wedge \text{eval}_{\text{Disk}^\xi}(s') = M \\ M & s = s' \cdot \alpha:\text{read}(b) \cdot \alpha:v \wedge \text{eval}_{\text{Disk}^\xi}(s') = M \wedge M[b] = v \\ \perp & \text{otherwise} \end{cases} \\
 \nu_{\text{Disk}^\xi} &:= \{s \mid \exists s'. s \sqsubseteq s' \wedge \text{eval}_{\text{Disk}^\xi}(s') \neq \perp\}
 \end{aligned}$$

```

MFS:
Import disk:Disk
Import lockmap:LockMapB
Import log:Log

void write(block_id dir, file_id fid, data_block data) {
    lock[dir].acquire();
    dir_inode ← (folder_inode) disk.read(dir);
    file ← dir_inode[fid];
    disk.write(file, data);
    lock[dir].release();
}

data_block read(block_id dir, file_id fid) {
    lock[dir].acquire();
    dir_inode ← (folder_inode) disk.read(dir);
    file ← dir_inode[fid];
    data ← disk.read(file);
    lock[dir].release();
    return data;
}

void swap(block_id src, block_id tgt, file_id src_f, file_id tgt_f) {
    if (src < tgt) {
        lock[src].acquire();
        lock[tgt].acquire();
    } else {
        lock[tgt].acquire();
        lock[src].acquire();
    }
    src_inode ← (folder_inode) disk.read(src);
    tgt_inode ← (folder_inode) disk.read(tgt);
    src_file ← src_inode[src_f];
    tgt_file ← tgt_inode[tgt_f];
    log.insert(src, tgt, src_f, tgt_f, src_file, tgt_file);

    disk.write(tgt, tgt_inode[tgt_f -> src_file]);
    disk.write(src, src_inode[src_f -> tgt_file]);
    log.remove( $\alpha$ );
    lock[src].release();
    lock[tgt].release();
}

void recovery() {
    for (i = 0; i < |agents|; i++) {
        entry ← log.get(agents[i]);
        if (entry = None)
            continue;
        (src, tgt, src_f, tgt_f, src_file, tgt_file) ← entry;

        src_inode ← (folder_inode) disk.read(src);
        tgt_inode ← (folder_inode) disk.read(tgt);

        disk.write(tgt, tgt_inode[tgt_f -> src_file]);
        disk.write(src, src_inode[src_f -> tgt_file]);
        log.remove(agents[i]);
    }
}

```

Fig. 22. Implementation of the File System

The LockMapB object is specified as a collection of individual locks, indexed by block_id type.

$$\begin{aligned}
\text{LockMapB} &:= \{\text{acq} : \text{block_id} \rightarrow 1, \text{rel} : \text{block_id} \rightarrow 1\} \\
\text{eval}_{\text{LockMapB}} &: P_{\dagger\text{LockMapB}} \rightarrow \{\perp\} + (\text{block_id} \xrightarrow{\text{fin}} \mathcal{P}(\Upsilon)) \\
\text{eval}_{\text{LockMapB}}(s) &:= \begin{cases} \bigcup_{b \in \text{block_id}} \{f \leftarrow \emptyset\} & \epsilon \\ m[b \mapsto \{\alpha\}] & s = s' \cdot \alpha:\text{acq}(b) \cdot \alpha:\text{ok} \wedge \text{eval}_{\text{LockMapB}}(s') = m \wedge m[b] = \emptyset \\ m[b \mapsto \emptyset] & s = s' \cdot \alpha:\text{rel}(b) \cdot \alpha:\text{ok} \wedge \text{eval}_{\text{LockMapB}}(s') = m \wedge m[b] = \{\alpha\} \\ \perp & \text{otherwise} \end{cases} \\
\nu_{\text{LockMapB}^\sharp} &:= \text{vol}(\{s \mid \exists s'. s \sqsubseteq s' \wedge \text{eval}_{\text{LockMapB}}(s') \neq \perp\})
\end{aligned}$$

The lock map is in fact equivalent to a horizontal composition of volatile locks:

$$\nu_{\text{LockMapB}^\sharp} := \otimes_{b \in \text{block_id}} \text{vol}(\nu_{\text{Lock}})$$

Since Oliveira Vale et al. [31] have verified the linearizability of a lock (in fact, the ticket lock implementation M_{Lock}) we may lift their proof to our setting using Prop. 2.8, as explained in the main paper.

The last object Log is a write ahead log used for crash atomicity of swap operations. It stores at most one log entry per thread, and each log entry contains all the information about a single swap operation: the source and target folder block id, the source and target file id, and the block id of the swapped files. While a thread may only insert entry for itself, it can remove entry for any thread due to the recovery routine. Otherwise, the specification of Log is simply another map, with the following formal specification,

$$\begin{aligned}
\text{entry} &:= \text{block_id} \times \text{block_id} \times \text{file_id} \times \text{file_id} \times \text{block_id} \times \text{block_id} \\
\text{Log} &:= \{\text{insert} : \text{entry} \rightarrow 1, \text{get} : \Upsilon \rightarrow \text{option entry}, \text{remove} : \Upsilon \rightarrow 1\} \\
\text{eval}_{\text{Log}^\sharp} &: P_{\dagger\text{Log}^\sharp} \rightarrow \{\perp\} + (\Upsilon \xrightarrow{\text{fin}} \text{entry}) \\
\text{eval}_{\text{Log}^\sharp}(s) &:= \begin{cases} \bigcup_{\alpha \in \Upsilon} [\alpha \mapsto \text{None}] & \epsilon \\ l & s = s' \cdot \sharp \wedge \text{eval}_{\text{Log}^\sharp}(s') = l \\ l & s = s' \cdot m \cdot \sharp \wedge \text{eval}_{\text{Log}^\sharp}(s' \cdot \sharp) = l \wedge \lambda_{\text{Log}}(m) = O \\ l[\alpha \mapsto e] & s = s' \cdot \alpha:\text{insert}(e) \cdot \alpha:\text{ok} \wedge \text{eval}_{\text{Log}^\sharp}(s') = l \\ l & s = s' \cdot \alpha:\text{get}(\alpha') \cdot \alpha:l[\alpha'] \wedge \text{eval}_{\text{Log}^\sharp}(s') = l \\ l[\alpha' \mapsto \text{None}] & s = s' \cdot \alpha:\text{remove}(\alpha') \cdot \alpha:\text{ok} \wedge \text{eval}_{\text{Log}^\sharp}(s') = l \\ \perp & \text{otherwise} \end{cases} \\
\nu_{\text{Log}^\sharp} &:= \{s \mid \exists s'. s \sqsubseteq s' \wedge \text{eval}_{\text{Log}^\sharp}(s') \neq \perp\}
\end{aligned}$$

Finally, the specification of the file system FS is a nested map from block ids (of folders) into a map from file ids to file contents, and the formal definition is given below,

$$\begin{aligned}
\text{FS} &:= \left\{ \begin{array}{l} \text{write} : \text{block_id} \times \text{file_id} \times \text{block} \rightarrow 1, \\ \text{read} : \text{block_id} \times \text{file_id} \rightarrow \text{block}, \\ \text{swap} : \text{block_id} \times \text{block_id} \times \text{file_id} \times \text{file_id} \rightarrow 1 \end{array} \right\} \\
&\quad f[a \mapsto b \mapsto c] := f[a \mapsto [f[a][b \mapsto c]]] \\
\text{eval}_{\text{FS}^\sharp} &: P_{\dagger\text{FS}^\sharp} \rightarrow \{\perp\} + (\text{block_id} \xrightarrow{\text{fin}} \text{file_id} \xrightarrow{\text{fin}} \text{block}) \\
\text{eval}_{\text{FS}^\sharp}(s) &:= \begin{cases} \bigcup_{\alpha \in \Upsilon} [\alpha \mapsto \text{None}] & \epsilon \\ f & s = s' \cdot \underline{f} \wedge \text{eval}_{\text{FS}^\sharp}(s') = f \\ f & s = s' \cdot m \cdot \underline{f} \wedge \text{eval}_{\text{FS}^\sharp}(s' \cdot \underline{f}) = fs \wedge \lambda_{\text{FS}}(m) = O \\ f[a \mapsto b \mapsto c] & s = s' \cdot \alpha : \text{write}(a, b, c) \cdot \alpha : \text{ok} \wedge \text{eval}_{\text{FS}^\sharp}(s') = f \\ f & s = s' \cdot \alpha : \text{read}(a, b) \cdot \alpha : f[a][b] \wedge \text{eval}_{\text{FS}^\sharp}(s') = f \\ f[a \mapsto b \mapsto f[c][d]][c \mapsto d \mapsto f[a][b]] & s = s' \cdot \alpha : \text{swap}(a, b, c, d) \cdot \alpha : \text{ok} \wedge \text{eval}_{\text{FS}^\sharp}(s') = f \\ \perp & \text{otherwise} \end{cases} \\
v_{\text{FS}^\sharp} &:= \{s \mid \exists s'. s \sqsubseteq s' \wedge \text{eval}_{\text{FS}^\sharp}(s') \neq \perp\}
\end{aligned}$$

1.3.2 Proof of the File System. The line-by-line proof is presented in Figure 23, Figure 24, Figure 25, Figure 26, and we highlight the important steps here.

First, we define the maximal linearized prefix of current possibility ρ ,

$$\begin{aligned}
\text{lin} &: P_{\dagger\text{FS}^\sharp} \rightarrow P_{\dagger\text{FS}^\sharp} \\
\text{lin}(\rho) = p_{\underline{f}} \cdot p &\iff \left(\begin{array}{l} p_{\underline{f}} \cdot p \sqsubseteq \rho \wedge \\ (p_{\underline{f}} = \epsilon \vee \exists p'. p' \cdot \underline{f} = p_{\underline{f}}) \wedge \\ p \in P_{\dagger\text{FS}} \wedge \forall p'. p \sqsubseteq p' \wedge p_{\underline{f}} \cdot p' \sqsubseteq \rho \implies p' \notin P_{\dagger\text{FS}} \end{array} \right)
\end{aligned}$$

which helps us derive the current state of the object according to the possibility as well as the concrete play,

$$\begin{aligned}
\text{state}^\rho &: P_{\dagger\text{FS}^\sharp} \rightarrow \{\perp\} + (\text{block_id} \xrightarrow{\text{fin}} \text{file_id} \xrightarrow{\text{fin}} \text{block}) \\
\text{state}^\rho(\rho) &:= \text{eval}_{\text{FS}^\sharp}(\text{lin}(\rho)) \\
\text{state}^s &: P_{\dagger(\text{Disk}\&\text{Log}\&\text{LockMapB})^\sharp} \rightarrow \{\perp\} + (\text{file_id} \xrightarrow{\text{fin}} \text{block}) \\
\text{state}^s(s) &:= \text{eval}_{\text{Disk}^\sharp}(s \upharpoonright_{\text{Disk}^\sharp})
\end{aligned}$$

Since the directory inodes are never moved around, we take the liberty to use notation $\text{state}^s(s)[d][f]$ when it's clear that $d \in \text{block_id}$ is a folder block and f is a file id. Next, we are interested in the current status of lock ownership,

$$\begin{aligned}
\text{owned} &: \mathcal{P}(P_{\dagger(\text{Disk}\&\text{Log}\&\text{LockMapB})^\sharp} \times \text{block_id}) \\
\text{owned}(s, b) &\iff \text{eval}_{\text{LockMapB}^\sharp}(s \upharpoonright_{\text{LockMapB}^\sharp})[b] \neq \emptyset \\
\text{ownedby} &: P_{\dagger(\text{Disk}\&\text{Log}\&\text{LockMapB})^\sharp} \times \Upsilon \rightarrow \mathcal{P}(\text{block_id}) \\
\text{ownedby}(s, \alpha) &:= \{b \mid \text{eval}_{\text{LockMapB}^\sharp}(s \upharpoonright_{\text{LockMapB}^\sharp})[b] = \{\alpha\}\}
\end{aligned}$$

Finally, we care about whether certain blocks are currently mentioned by some entry in WAL,

$$\begin{aligned} \text{logged} &: \mathcal{P}(P_{\dagger}(\text{Disk\&Log\&LockMapB})^{\sharp} \times \text{block_id} \times \text{file_id}) \\ \text{logged}(s, d, f) &\iff \exists \alpha, d_1, d_2, f_1, f_2, b_1, b_2. \left(\begin{array}{l} \text{eval}_{\text{Log}^{\sharp}}(s \upharpoonright_{\text{Log}^{\sharp}})[\alpha] = (d_1, d_2, f_1, f_2, b_1, b_2) \wedge \\ ((d, f) = (d_1, f_1) \vee (d, f) = (d_2, f_2)) \end{array} \right) \\ \text{logged}^2 &: \mathcal{P}(P_{\dagger}(\text{Disk\&Log\&LockMapB})^{\sharp} \times \text{block_id} \times \text{block_id} \times \text{file_id} \times \text{file_id}) \\ \text{logged}^2(s, d_1, d_2, f_1, f_2) &\iff \exists \alpha, b_1, b_2. \left(\begin{array}{l} \text{eval}_{\text{Log}^{\sharp}}(s \upharpoonright_{\text{Log}^{\sharp}})[\alpha] = (d_1, d_2, f_1, f_2, b_1, b_2) \vee \\ \text{eval}_{\text{Log}^{\sharp}}(s \upharpoonright_{\text{Log}^{\sharp}})[\alpha] = (d_2, d_1, f_2, f_1, b_1, b_2) \end{array} \right) \end{aligned}$$

The runtime invariant of the program is then a collection of observations. Firstly, the file content we compute from the two different states matches with each other except for those currently in a WAL entry,

$$I_1(s, \rho) \iff \forall d, f. \neg \text{logged}(d, f) \implies \text{state}^s(s)[\text{state}^s(s)[d][f]] = \text{state}^{\rho}(\rho)[d][f]$$

Secondly, two file blocks can only be duplicates of each other only if both of them are in the same WAL entry,

$$I_2(s, \rho) \iff \forall d_1, d_2, f_1, f_2. \text{state}^s(s)[d_1][f_1] = \text{state}^s(s)[d_2][f_2] \implies \text{logged}^2(s, d_1, d_2, f_1, f_2)$$

Thirdly, all entries in WAL corresponds to a pending invocation (potentially before some crashes) in the possibility, and the recorded file block id matches the current overlay state (before swap),

$$I_3(s, \rho) \iff \forall \alpha, d_1, d_2, f_1, f_2, b_1, b_2. \left(\begin{array}{l} \text{eval}_{\text{Log}^{\sharp}}(s \upharpoonright_{\text{Log}^{\sharp}})[\alpha] = (d_1, d_2, f_1, f_2, b_1, b_2) \implies \\ \text{state}^{\rho}(\rho)[d_1][f_1] = b_1 \wedge \text{state}^{\rho}(\rho)[d_2][f_2] = b_2 \wedge \\ \exists p. \pi_{\alpha}(\rho) \upharpoonright_{\text{FS}} = p \cdot \text{swap}(d_1, d_2, f_1, f_2) \end{array} \right)$$

Finally, membership in WAL implies lock ownership by the same thread,

$$I_4(s, \rho) \iff \forall \alpha, d_1, d_2, f_1, f_2, b_1, b_2. \left(\begin{array}{l} \text{eval}_{\text{Log}^{\sharp}}(s \upharpoonright_{\text{Log}^{\sharp}})[\alpha] = (d_1, d_2, f_1, f_2, b_1, b_2) \implies \\ \{d_1, d_2\} = \text{ownedby}(s, \alpha) \end{array} \right)$$

The runtime invariant (both the precondition and postcondition of all FS methods) is then the conjunction of all above,

$$I := I_1 \cap I_2 \cap I_3 \cap I_4$$

However, due to the fact that the locks are not persistent, the last conjunction is not stable with respect to crashes. Thus the crash invariant is the conjunction of the first three,

$$\begin{aligned} I_{\sharp} &:= I_1 \cap I_2 \cap I_3 \\ (s, \rho) Q_{\sharp}(s', \rho') &\iff I_{\sharp}(s', \rho') \end{aligned}$$

The precondition and postcondition of any method (except for recovery) assumes the invariant as well as empty ownership of current thread,

$$\begin{aligned} P^f(\Delta, s, \rho) &\iff I(s, \rho) \wedge \text{ownedby}(s, \alpha) = \emptyset \\ (-) Q^f(\Delta, s', \rho') &\iff I(s', \rho') \wedge \text{ownedby}(s', \alpha) = \emptyset \end{aligned}$$

The guarantee condition, in addition to preservation of invariants, further specifies that one thread may only update folders or files they currently owns, physically or abstractly. Rely condition

is then the union of guarantee and method invocations and returns,

$$(s, \rho) \mathcal{G}[\alpha] (s', \rho') \iff \left(\begin{array}{l} \forall d. d \notin \text{ownedby}(s, \alpha) \implies \text{state}^s(s)[b] = \text{state}^s(s')[b] \wedge \\ \forall d. d \notin \text{ownedby}(s, \alpha) \implies \text{state}^\rho(s)[d] = \text{state}^\rho(s')[d] \wedge \\ I(s', \rho') \end{array} \right)$$

$$\mathcal{R}[\alpha] \triangleq \bigcup_{\alpha' \in Y, \alpha' \neq \alpha} \mathcal{G}[\alpha'] \cup \text{invoke}_{\alpha'}(-) \cup \text{return}_{\alpha'}(-)$$

The linearization point of swap normally happens at the time when log entry is removed, which by itself is a straightforward proof thanks to mutual exclusion. However, if a swap method crashes after inserting the WAL entry but before removing the entry, it will instead be retroactively linearized during the recovery procedure, also at the point when the log is removed from WAL. While the same operation may be performed multiple times and some are even partially performed, it is safe nonetheless thanks to the idempotent nature of log entry application.

The linearization point of write happens at the point when the underlay actually writes to the disk and read linearizes at the disk read operation, as one would expect. The safety is provided by the mutual exclusion of locks, the same as the swap operation.

```

    {I(s, ρ) ∧ ownedby(s, α) = ∅}
1: write(dir, fid, data) {
2:   lock[dir].acquire();
   {I(s, ρ) ∧ ownedby(s, α) = {dir}}
3:   dir_inode ← (folder_inode)disk.read(dir);
4:   file ← dir_inode[fid];
   {I(s, ρ) ∧ ownedby(s, α) = {dir} ∧ states(s)[dir][fid] = file}
5:   disk.write(file, data);
   {I(s, ρ) ∧ ownedby(s, α) = {dir} ∧ states(s)[dir][fid] = file ∧ states(s)[file] = data}

   { I(s', ρ') ∧ (
     ownedby(s', α) = {dir} ∧ states(s')[dir][fid] = file ∧ states(s')[file] = data ∧
     lin(ρ) · α:write(dir, fid, data) · α:ok ⊆ ρ'
   ) }
6:   lock[dir].release();
   {I(s, ρ) ∧ ownedby(s, α) = ∅ ∧ ∃p. πα(ρ) = p · ok}
7: }
{I(s, ρ) ∧ ownedby(s, α) = ∅ ∧ returned[write](Δ, s, ρ)}

```

Fig. 23. Proof of file system - write

1.3.3 Replicated Disk. To demonstrate vertical composition, we implement the aforementioned disk interface on top of several disks. Thanks to the locality property of tensor operator, we only need to verify a single disk block and safely compose into the whole disk with little effort. The

```

    {I(s, ρ) ∧ ownedby(s, α) = ∅}
1: write(dir, fid){
2:   lock[dir].acquire();
   {I(s, ρ) ∧ ownedby(s, α) = {dir}}
3:   dir_inode ← (folder_inode)disk.read(dir);
4:   file ← dir_inode[fid];
   {I(s, ρ) ∧ ownedby(s, α) = {dir} ∧ states(s)[dir][fid] = file}
5:   data ← disk.read(file);
   {I(s, ρ) ∧ ownedby(s, α) = {dir} ∧ states(s)[dir][fid] = file ∧ states(s)[file] = data}

   { I(s', ρ') ∧ (
     ownedby(s', α) = {dir} ∧ states(s')[dir][fid] = file ∧ states(s')[file] = data ∧
     lin(ρ) · α:read(data) · α:data ⊆ ρ'
   ) }
6:   lock[dir].release();
   {I(s, ρ) ∧ ownedby(s, α) = ∅ ∧ ∃p.πα(ρ) = p · data}

7: }
    {I(s, ρ) ∧ ownedby(s, α) = ∅ ∧ returned[write](Δ, s, ρ)}
    
```

Fig. 24. Proof of file system - read

single block specification, which both the underlay and overlay follows, is provided below,

$$\begin{aligned}
 \text{DiskBlock} &:= \{\text{write} : \text{block} \rightarrow 1, \text{read} : 1 \rightarrow \text{block}\} \\
 \text{eval}_{\text{DiskBlock}^\xi} &: P_{\dagger\text{DiskBlock}^\xi} \rightarrow \{\perp\} + \text{block} \\
 \text{eval}_{\text{DiskBlock}^\xi}(s) &:= \begin{cases} 0 & \epsilon \\ v & s = s' \cdot \zeta \wedge \text{eval}_{\text{Disk}^\xi}(s') = v \\ v & s = s' \cdot m \cdot \zeta \wedge \text{eval}_{\text{Disk}^\xi}(s' \cdot \zeta) = v \wedge \lambda_{\text{Disk}}(m) = O \\ v & s = s' \cdot \alpha:\text{write}(v) \cdot \alpha:\text{ok} \wedge \text{eval}_{\text{Disk}^\xi}(s') \neq \perp \\ v & s = s' \cdot \alpha:\text{read}() \cdot \alpha:v \wedge \text{eval}_{\text{Disk}^\xi}(s') = v \\ \perp & \text{otherwise} \end{cases} \\
 v_{\text{DiskBlock}^\xi} &:= \{s \mid \exists s'.s \sqsubseteq s' \wedge \text{eval}_{\text{DiskBlock}^\xi}(s') \neq \perp\}
 \end{aligned}$$

The implementation is given in Figure 27. In the implementaion, we do not acquire locks when writing or reading disk blocks. This is sound since the file system always guarantees mutual exclusion when calling disk operations on the same block. We express this fact through a client policy that requires the client to only access the disk atomically:

$$\xi_{\text{DiskBlock}} := P!_{\text{DiskBlock}}$$

where $P!_{\text{DiskBlock}}$ is the set of well-formed atomic plays over DiskBlock .

1.3.4 Verification of Replicated Disks. For the purpose of verification, we assume the lin function and the eval function are defined in the same way as in the file interface verification. Following a similar approach, we define state^ρ and state^s according to the possibility and the the underlay play respectively. We use $\text{DiskBlock}[N]_i$ to denote the i -th disk of the N disk array of the underlay.

$$\begin{aligned}
 \text{state}^\rho &: P_{\dagger\text{DiskBlock}^\xi} \rightarrow \{\perp\} + \text{block} \\
 \text{state}^\rho(\rho) &:= \text{eval}_{\text{DiskBlock}^\xi}(\text{lin}(\rho)) \\
 \text{state}^s &: P_{\dagger\text{DiskBlock}[N]^\xi} \rightarrow \{\perp\} + \text{block} \\
 \text{state}^s(\rho) &:= \text{eval}_{\text{DiskBlock}^\xi}(s \upharpoonright_{\text{DiskBlock}[N]_0^\xi})
 \end{aligned}$$

```

    {I(s, ρ) ∧ ownedby(s, α) = ∅}
1: swap(src, tgt, src_f, tgt_f) {
2:   if(src < tgt) {
3:     lock[src].acquire();
4:     lock[tgt].acquire();
5:   } else {
6:     lock[tgt].acquire();
7:     lock[src].acquire();
8:   }
   {I(s, ρ) ∧ ownedby(s, α) = {src, tgt}}
9:   src_inode ← (folder_inode)disk.read(src);
10:  tgt_inode ← (folder_inode)disk.read(tgt);
11:  src_file ← src_inode[src_f];
12:  tgt_file ← tgt_inode[tgt_f];
   {I(s, ρ) ∧ (
   ownedby(s, α) = {src, tgt} ∧
   states(s)[src] = src_inode ∧ states(s)[tgt] = tgt_inode)}
13:  log.insert(src, tgt, src_f, tgt_f, src_file, tgt_file);
   {I(s, ρ) ∧ (
   ownedby(s, α) = {src, tgt} ∧
   logged2(s, src, tgt, src_f, tgt_f) ∧
   states(s)[src] = src_inode ∧ states(s)[tgt] = tgt_inode)}
14:  disk.write(tgt, tgt_inode[tgt_f ↦ src_file]);
   {I(s, ρ) ∧ (
   ownedby(s, α) = {src, tgt} ∧
   logged2(s, src, tgt, src_f, tgt_f) ∧
   states(s)[src] = src_inode ∧ states(s)[tgt] = tgt_inode[tgt_f ↦ src_file])}
15:  disk.write(src, src_inode[src_f ↦ tgt_file]);
   {I(s, ρ) ∧ (
   ownedby(s, α) = {src, tgt} ∧
   logged2(s, src, tgt, src_f, tgt_f) ∧
   states(s)[src] = src_inode[src_f ↦ tgt_file] ∧ states(s)[tgt] = tgt_inode[tgt_f ↦ src_file])}
16:  log.remove(α);
   {I(s, ρ) ∧ (
   ownedby(s, α) = {src, tgt} ∧
   states(s)[src] = src_inode[src_f ↦ tgt_file] ∧ states(s)[tgt] = tgt_inode[tgt_f ↦ src_file])}

   {I(s', ρ') ∧ (
   lin(ρ) · α:swap(src, tgt, src_f, tgt_f) · α:ok ⊆ ρ' ∧
   ownedby(s', α) = {src, tgt})}
   {I(s, ρ) ∧ ownedby(s, α) = {src, tgt} ∧ ∃p.πα(ρ) = p · ok}
17:  lock[src].release();
18:  lock[tgt].release();
   {I(s, ρ) ∧ ownedby(s, α) = ∅ ∧ ∃p.πα(ρ) = p · ok}
19: }
   {I(s, ρ) ∧ ownedby(s, α) = ∅ ∧ returned[swap](Δ, s, ρ)}

```

Fig. 25. Proof of file system - swap

We first give the invariant of this object,

$$I(s, \rho) \iff \text{state}^\rho(\rho) = \text{state}^s(s)$$

We further require that at the boundary of overlay methods, all disk blocks match the content of the possibility,

$$P^f(\Delta, s, \rho) \iff I(s, \rho) \wedge \forall i. 0 \leq i < N \implies \text{eval}_{\text{DiskBlock}^i}(s \upharpoonright_{\text{DiskBlock}[N]_i^i}) = \text{state}^\rho(\rho)$$

$$(-) Q^f(\Delta, s', \rho') \iff I(s', \rho') \wedge \forall i. 0 \leq i < N \implies \text{eval}_{\text{DiskBlock}^i}(s' \upharpoonright_{\text{DiskBlock}[N]_i^i}) = \text{state}^\rho(\rho')$$

```

    {Iz(s, ρ)}
1: recovery() {
2:   for(i = 0; i < |agents|; i++) {
      {Iz(s, ρ) ∧ ∀j.0 ≤ j < i ⇒ evalLogz(s↑Logz)[agents[j]] = None}
3:   entry ← log.get(agents[i]);
4:   if(entry = None)
5:     continue;
6:   (src, tgt, src_f, tgt_f, src_file, tgt_file) ← entry;
      {Iz(s, ρ) ∧ (
        ∀j.0 ≤ j < i ⇒ evalLogz(s↑Logz)[agents[j]] = None ∧
        (src, tgt, src_f, tgt_f, src_file, tgt_file) = evalLogz(s↑Logz)[agents[i]] ∧
        ∃p.πα(ρ) ↓Log = p · swap(src, tgt, src_f, tgt_f)
      )}
7:   src_inode ← (folder_inode)disk.read(src);
8:   tgt_inode ← (folder_inode)disk.read(tgt);
      {Iz(s, ρ) ∧ (
        ∀j.0 ≤ j < i ⇒ evalLogz(s↑Logz)[agents[j]] = None ∧
        (src, tgt, src_f, tgt_f, src_file, tgt_file) = evalLogz(s↑Logz)[agents[i]] ∧
        ∃p.πα(ρ) ↓Log = p · swap(src, tgt, src_f, tgt_f) ∧
        states(s)[src] = src_inode ∧ states(s)[tgt] = tgt_inode
      )}
9:   disk.write(tgt, tgt_inode[tgt_f ↦ src_file]);
      {Iz(s, ρ) ∧ (
        ∀j.0 ≤ j < i ⇒ evalLogz(s↑Logz)[agents[j]] = None ∧
        (src, tgt, src_f, tgt_f, src_file, tgt_file) = evalLogz(s↑Logz)[agents[i]] ∧
        ∃p.πα(ρ) ↓Log = p · swap(src, tgt, src_f, tgt_f) ∧
        states(s)[src] = src_inode ∧ states(s)[tgt] = tgt_inode[tgt_f ↦ src_file])
      )}
10:  disk.write(src, src_inode[src_f ↦ tgt_file]);
      {Iz(s, ρ) ∧ (
        ∀j.0 ≤ j < i ⇒ evalLogz(s↑Logz)[agents[j]] = None ∧
        (src, tgt, src_f, tgt_f, src_file, tgt_file) = evalLogz(s↑Logz)[agents[i]] ∧
        ∃p.πα(ρ) ↓Log = p · swap(src, tgt, src_f, tgt_f) ∧
        states(s)[src] = src_inode[src_f ↦ tgt_file] ∧ states(s)[tgt] = tgt_inode[tgt_f ↦ src_file])
      )}
11:  log.remove(agents[i]);
      {Iz(s, ρ) ∧ (
        ∀j.0 ≤ j ≤ i ⇒ evalLogz(s↑Logz)[agents[j]] = None ∧
        ∃p.πα(ρ) ↓Log = p · swap(src, tgt, src_f, tgt_f) ∧
        states(s)[src] = src_inode[src_f ↦ tgt_file] ∧ states(s)[tgt] = tgt_inode[tgt_f ↦ src_file])
      )}

      {Iz(s', ρ') ∧ (
        ∀j.0 ≤ j ≤ i ⇒ evalLogz(s'↑Logz)[agents[j]] = None ∧
        lin(ρ) · α:swap(src, tgt, src_f, tgt_f) · α:ok ⊆ ρ'
      )}
12: }
    {Iz(s, ρ) ∧ ∀α ∈ agents.evalLogz(s↑Logz)[α] = None}
13: }
    {I(s, ρ)}

```

Fig. 26. Proof of file system - recovery

We can now define the crash-postcondition as preservation of invariant,

$$(s, \rho) Q_z(s', \rho') \iff I(s', \rho')$$

since the invariant is preserved, the sole purpose of the recovery method is to re-establish the universal precondition of normal routines.

```

MDiskBlock:
Import blocks: DiskBlock[N]

void write(block_data data) {
  for (i ← 0; i < N; i ← i + 1) {
    blocks[i].write(data);
  }
}

block_data read() {
  i ← random();
  block_data data ← blocks[i].read();
  return data;
}

void recover() {
  data ← disks[0].read();
  for (i ← 0; i < N; i ← i + 1) {
    disks[i].write(data);
  }
}

```

Fig. 27. Implementation for Replicated Disks

The rely and guarantee condition is trivial since the client policy effectively disallow any type of interleaving,

$$(s, \rho) \mathcal{G}[\alpha] (s', \rho') \iff \text{id}$$

$$\mathcal{R}[\alpha] \triangleq \bigcup_{\alpha' \in \Upsilon, \alpha' \neq \alpha} \mathcal{G}[\alpha'] \cup \text{invoke}_{\alpha'}(-) \cup \text{return}_{\alpha'}(-)$$

With everything defined, the step through proofs are given in Figure 28, Figure 29, and Figure 30. For brevity, we use $\text{eval}_{\text{DiskBlock}[N]_i^\xi}(s)$ as a abbreviation for $\text{eval}_{\text{Disk}^\xi}(s \upharpoonright_{\text{DiskBlock}[N]_i})$.

```

{I(s, ρ)}
1: write(data){
2:  disks[0].write(data);
   {I(s, ρ) ∧ states(s) = stateρ(ρ) = data}
   {
     I(s', ρ') ∧ (states(s') = stateρ(ρ') = data ∧
     lin(ρ) · α:write(data) · α:ok ⊆ ρ')
   }
   {I(s, ρ) ∧ (∃p.πα(ρ) = p · ok) ∧ states(s) = stateρ(ρ) = data}
3:  for(i ← 1; i < N; i ← i + 1)
   {
     {I(s, ρ) ∧ (∃p.πα(ρ) = p · ok) ∧ ∀i'.0 ≤ i' < i ⇒ evalDiskBlock[N]_i'ξ(s) = stateρ(ρ) = data}
4:     disks[i].write(data);
     {
       {I(s, ρ) ∧ (∃p.πα(ρ) = p · ok) ∧ ∀i'.0 ≤ i' ≤ i ⇒ evalDiskBlock[N]_i'ξ(s) = stateρ(ρ) = data}
       {I(s, ρ) ∧ (∃p.πα(ρ) = p · ok) ∧ ∀i'.0 ≤ i' ≤ N ⇒ evalDiskBlock[N]_i'ξ(s) = stateρ(ρ) = data}
     }
5:   }
   {I(s, ρ) ∧ returned[write](Δ, s, ρ)}

```

Fig. 28. Proof of Replicated Disks - write

```

    {I(s, ρ)}
1: read(){
2:   i ← random(0, N - 1);
   {I(s, ρ) ∧ 0 ≤ i < N}
3:   data ← disks[i].read();
   {I(s, ρ) ∧ 0 ≤ i < N ∧ stateρ(ρ) = data}

   {I(s', ρ') ∧ stateρ(ρ')[b] = data ∧ lin(ρ) · α:read · α:data ⊆ ρ'}
   {I(s, ρ) ∧ (∃p.πα(ρ) = p · data)}
4:   return data;
5: }
    {I(s, ρ) ∧ returned[read](Δ, s, ρ)}

```

Fig. 29. Proof of Replicated Disks - read

```

    {I(s, ρ)}
1: recover(){
2:   data ← disks[0].read();
   {I(s, ρ) ∧ data = stateρ(ρ)}
3:   for(i ← 1; i < N; i ← i + 1)
     {I(s, ρ) ∧ ∀i'.0 ≤ i' < i ⇒ evalDiskBlock[N]ii'(s) = stateρ(ρ) = data}
4:   disks[i].write(b, data);
     {I(s, ρ) ∧ ∀i'.0 ≤ i' ≤ i ⇒ evalDiskBlock[N]ii'(s) = stateρ(ρ) = data}
     {I(s, ρ) ∧ ∀i'.0 ≤ i' < N ⇒ evalDiskBlock[N]ii'(s) = stateρ(ρ) = data}
5: }
    {I(s, ρ) ∧ ∀i'.0 ≤ i' < N ⇒ evalDiskBlock[N]ii'(s) = stateρ(ρ) = data}

```

Fig. 30. Proof of Replicated Disks - recovery

1.3.5 Write-Ahead Log Implementation. The implementation of the write-ahead log used in the file system is presented in Figure 31. We omit the verification as it is straightforward in terms of crash linearizability: all operations are immediately persisted and there is no in-between states for the disk block.

Because the disk is equivalent to the horizontal composition of its blocks:

$$V_{\text{Disk}^i} \cong \otimes_{b \in \text{block_id}} V_{\text{DiskBlock}^i}$$

It follows that we may separate one location from the disk for the log, as follows:

$$\begin{aligned}
 & (\text{vol}(V_{\text{Buffer}}) \otimes \text{vol}(V_{\text{Lock}})) \otimes V_{\text{Disk}^i} \\
 & \cong (\text{vol}(V_{\text{Buffer}}) \otimes \text{vol}(V_{\text{Lock}})) \otimes (\otimes_{b \in \text{block_id}} V_{\text{DiskBlock}^i}) \\
 & \cong (V_{\text{DiskBlock}^i} \otimes \text{vol}(V_{\text{Buffer}}) \otimes \text{vol}(V_{\text{Lock}})) \otimes \otimes_{b \in \text{block_id} \setminus \text{logblk}} V_{\text{DiskBlock}^i}
 \end{aligned}$$

so that

$$(\text{vol}(V_{\text{Buffer}}) \otimes \text{vol}(V_{\text{Lock}}) \otimes V_{\text{Disk}^i}); (M_{\text{Log}} \otimes \text{crashcopy})$$

is linearizable to $V_{\text{Log}^i} \otimes (\otimes_{b \in \text{block_id} \setminus \text{logblk}} V_{\text{DiskBlock}^i})$ where `logblk` is the block id where the log is located. That is to say, we obtain a log together with a disk with size one block less than before, such that in the underlay the log lives in the same disk.

```

MLog:
Import block:DiskBlock
Import buffer:Buffer
Import lock:Lock

void insert(block_id a, block_id b, file_id c, file_id d, block_id e, block_id f) {
    lock.acquire();
    buffer[α] = (a, b, c, d, e, f);
    block.write(buffer);
    lock.release();
}

void remove(agent_id agent) {
    lock.acquire();
    buffer[agent] = None;
    block.write(buffer);
    lock.release();
}

void get(agent_id agent) {
    return buffer[agent];
}

void recovery() {
    buffer = (log) block.read();
}

```

Fig. 31. Implementation of the Write Ahead Log

J PROOFS

J.1 Basic Semicategorical Structure

In the following we assume that the operation $-; -$ is defined arbitrary sets of plays, instead of just on strategies. The definition is exactly the same. Given $s \in P_{A \rightarrow B}$ and $t \in P_{B \rightarrow C}$ we write $s; t$ for $\{s\}; \{t\}$. We also take the convention of writing, for $s \in P_A$ with $A \in \underline{\text{Crash}}$,

$$s = s_1 \cdot \frac{!}{!} \cdot s_2 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot s_{n+1}$$

where for each i , $s_i = \text{epo}_i(s)$.

The following important lemma characterizes composition in Crash using composition in Conc.

LEMMA J.1. *Given strategies $\sigma : A \multimap B$, $\tau : B \multimap C \in \underline{\text{Crash}}$, for any play $u \in \sigma; \tau$, there exists $s \in \sigma$, $t \in \tau$ such that*

$$s = s_1 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot s_{n+1} \quad t = t_1 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot t_{n+1}$$

and, for all i ,

$$s_i \in \mathbb{P}_{A^i \multimap B^i}^{\text{conc}} \quad \text{and} \quad t_i \in \mathbb{P}_{B^i \multimap C^i}^{\text{conc}}$$

and such that u decomposes as

$$u = u_1 \cdot \frac{!}{!} \cdot u_2 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot u_{n+1}$$

with $u_i \in s_i; t_i$ for all i .

PROOF. Suppose $u \in \sigma; \tau$. By definition there exists $u' \in \text{int}(\sigma, \tau)$ such that $u' \upharpoonright_{A, B, -} \in \sigma$, $u' \upharpoonright_{-, B, C} \in \tau$ and $u' \upharpoonright_{A, -, C} = u$, we claim that assigning $s := u' \upharpoonright_{A, B, -}$ and $t := u' \upharpoonright_{-, B, C}$ the claim is proven.

Since all three of u' , s , and t are well-formed it follows that they can be written as (it will soon be clear why all plays feature the same number of epochs):

$$s = s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1} \quad t = t_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot t_{n+1} \quad u' = u'_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot u'_{n+1}$$

So let $u_i = u'_i \upharpoonright_{A,-,C}$. By definition we know $u' \upharpoonright_{-,B,-} = s \upharpoonright_{-,B} = t \upharpoonright_{-,B}$, so it follows that the i -th crash signal in s matches t with the i -th crash signal in u' , and that they have the same number of crash signals. Furthermore, since $s = u' \upharpoonright_{A,B,-}$, the i -th epoch of u' , i.e. u'_i , projects to the i -th epoch of s , i.e. $u'_i \upharpoonright_{A,B,-} = s_i$. Similarly, $u'_i \upharpoonright_{-,B,C} = t_i$. Hence, it follows that $u'_i \upharpoonright_{A,-,C} = u_i \in s_i; t_i$. \square

PROPOSITION J.2. *strategy composition is well-defined and associative.*

PROOF. Well-defined Suppose $\sigma : A \multimap B$ and $\tau : B \multimap C$, then we have $\epsilon \in \sigma$ and $\epsilon \in \tau$. Then taking

$$\epsilon \in \text{int}(A, B, C)$$

we have $\epsilon \upharpoonright_{A,-,C} = \epsilon \in \mathbb{P}_{A \multimap C}^{\downarrow}$ which implies $\epsilon \in \sigma; \tau$. So $\sigma; \tau$ is non-empty

Now suppose $s \in \sigma; \tau$ and that $p \sqsubseteq s$, then there exists $s' \in \text{int}(\sigma, \tau)$ such that $s' \upharpoonright_{A,-,C} = s$. In particular $p \sqsubseteq s' \upharpoonright_{A,-,C}$. Hence there exists $p' \sqsubseteq s'$ such that $p' \upharpoonright_{A,-,C} = p$. Since $s' \upharpoonright_{A,B,-} \in \sigma$ and $s' \upharpoonright_{-,B,C} \in \tau$ and both σ, τ are prefix-closed, so $p' \upharpoonright_{A,B,-} \in \sigma$ and $p' \upharpoonright_{-,B,C} \in \tau$. So $p' \in \text{int}(\sigma, \tau)$. Since $\mathbb{P}_{A \multimap C}^{\downarrow}$ is prefix-closed, $p' \upharpoonright_{A,-,C} \sqsubseteq s' \upharpoonright_{A,-,C} \in \mathbb{P}_{A \multimap C}^{\downarrow}$ so $p' \upharpoonright_{A,-,C} \in \mathbb{P}_{A \multimap C}^{\downarrow}$. Hence $p \in \sigma; \tau$.

Lastly, we show $\sigma; \tau$ is \downarrow -receptive. Suppose $s \in \sigma; \tau$. So there exists $s' \in \text{int}(\sigma, \tau)$ such that $s' \upharpoonright_{A,-,C} = s$. Note then that we have $s' \cdot \downarrow \in \text{int}(A, B, C)$. By \downarrow -receptivity, it holds that $s' \cdot \downarrow \upharpoonright_{A,B,-} \in \sigma$ and $s' \cdot \downarrow \upharpoonright_{-,B,C} \in \tau$. So it follows that $s' \cdot \downarrow \upharpoonright_{A,-,C} = s \cdot \downarrow \in \sigma; \tau$.

Associative Suppose $\sigma : A \multimap B, \tau : B \multimap C$ and $\rho : C \multimap D$, fix $s \upharpoonright_{A,-,D} \in (\sigma; \tau); \rho$ where $u \in \text{int}((\sigma; \tau), \rho)$. Applying J.1 between the composition of $\sigma; \tau$ and ρ , and then applying the lemma again to decompose its projection to $\sigma; \tau$, we obtain that $u \upharpoonright_{A,-,D}$ can be written as

$$u_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot u_{n+1}$$

and $u'_i \in (s_i; t_i); r_i$ where $s_i \in \mathbb{P}_{A^Y \multimap B^Y}^{\text{conc}}$, $t_i \in \mathbb{P}_{B^Y \multimap C^Y}^{\text{conc}}$ and $r_i \in \mathbb{P}_{C^Y \multimap D^Y}^{\text{conc}}$, with, furthermore

$$s := s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1} \in \sigma$$

$$t := t_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot t_{n+1} \in \tau$$

$$r := r_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot r_{n+1} \in \rho$$

From concurrent games [31] we already know that $(s_i; t_i); r_i = s_i; (t_i; r_i)$. So there exists $u'_i \in \text{int}(s_i, t_i; r_i)$ such that $u'_i \upharpoonright_{A,-,D} = u_i$, $u'_i \upharpoonright_{A,B,-} = s_i$ and $u'_i \upharpoonright_{-,B,D} \in t_i; r_i$. Now let's define

$$u'' := u'_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot u'_{n+1}$$

Now we have $u'' \upharpoonright_{A,B,-} = s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1}$ and $u'' \upharpoonright_{-,B,D} \in t; r \subseteq \tau; \rho$. So $u = u'' \upharpoonright_{A,-,D} \in \sigma; (\tau; \rho)$.

The other direction is analogous. \square

We also prove that composition is monotonic and join-preserving, the main requirement to obtain an enriched semicategory.

PROPOSITION J.3. *For strategies*

$$\sigma : A \multimap B, \quad \tau : B \multimap C$$

the following hold:

(1) *if $\sigma \subseteq \sigma' : A \multimap B$ and $\tau \subseteq \tau' : B \multimap C$ then $\sigma; \tau \subseteq \sigma'; \tau'$*

(2) Given a family of strategies $(\sigma_i : A \multimap B)_{i \in I}$ it holds that

$$\left(\bigcup_{i \in I} \sigma_i \right); \tau = \bigcup_{i \in I} (\sigma_i; \tau)$$

(3) Given a family of strategies $(\tau_i : B \multimap C)_{i \in I}$ it holds that

$$\sigma; \left(\bigcup_{i \in I} \tau_i \right) = \bigcup_{i \in I} (\sigma; \tau_i)$$

PROOF. (1) Suppose $s \upharpoonright_{A,-,C} \in \sigma; \tau$ then

$$s \upharpoonright_{A,B,-} \in \sigma \implies s \upharpoonright_{A,B,-} \in \sigma'$$

$$s \upharpoonright_{-,B,C} \in \tau \implies s \upharpoonright_{-,B,C} \in \tau'$$

Also since $s \in \text{int}(A, B, C)$, $s \upharpoonright_{A,-,C} \in \mathbb{P}_{A \multimap C}^{\neq}$ so $s \upharpoonright_{A,-,C} \in \sigma'; \tau'$.

(2) For one direction, since we have $\sigma_i \subseteq \bigcup_{i \in I} \sigma_i$ so

$$\sigma_i; \tau \subseteq \left(\bigcup_{i \in I} \sigma_i \right); \tau$$

hence

$$\bigcup_{i \in I} (\sigma_i; \tau) \subseteq \left(\bigcup_{i \in I} \sigma_i \right); \tau$$

For the other direction, suppose $s \upharpoonright_{A,-,C} \in \left(\bigcup_{i \in I} \sigma_i \right); \tau$ which means $s \upharpoonright_{A,B,-} \in \left(\bigcup_{i \in I} \sigma_i \right)$ and $s \upharpoonright_{-,B,C} \in \tau$.

so there exists i such that $s \upharpoonright_{A,B,-} \in \sigma_i$, and therefore $s \upharpoonright_{A,-,C} \in \sigma_i; \tau \subseteq \bigcup_{i \in I} (\sigma_i; \tau)$ so that

$$s \upharpoonright_{A,-,C} \in \bigcup_{i \in I} (\sigma_i; \tau)$$

(3) For one direction we have $\tau_i \subseteq \bigcup_{i \in I} \tau_i$ so

$$\sigma; \tau_i \subseteq \sigma; \left(\bigcup_{i \in I} \tau_i \right)$$

hence

$$\bigcup_{i \in I} (\sigma; \tau_i) \subseteq \sigma; \left(\bigcup_{i \in I} \tau_i \right)$$

For the other direction, suppose $s \upharpoonright_{A,-,C} \in \sigma; \left(\bigcup_{i \in I} \tau_i \right)$ which means $s \upharpoonright_{A,B,-} \in \sigma$ and $s \upharpoonright_{-,B,C} \in \bigcup_{i \in I} \tau_i$.

so there exists i such that $s \upharpoonright_{-,B,C} \in \tau_i$, and therefore $s \upharpoonright_{A,-,C} \in \sigma; \tau_i \subseteq \bigcup_{i \in I} (\sigma; \tau_i)$ so that

$$s \upharpoonright_{A,-,C} \in \bigcup_{i \in I} (\sigma; \tau_i)$$

□

J.2 Volatile Lift and Idempotence

PROPOSITION J.4.

$$\text{vol}(-) : \underline{\text{Conc}} \rightarrow \underline{\text{Crash}}$$

is a semi-functor.

PROOF. For given

$$\sigma : A \multimap B \in \underline{\text{Conc}} \qquad \sigma' : B \multimap C \in \underline{\text{Conc}}$$

we want to show

$$\text{vol}(\sigma); \text{vol}(\sigma') = \text{vol}(\sigma; \sigma')$$

For one direction, fix $s \in \text{vol}(\sigma); \text{vol}(\sigma')$. By lemma J.1, we have that s can be decomposed as

$$s := s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1}$$

where for each i , $s_i \in \sigma; \sigma'$. It then follows immediately by the definition of $\text{vol}(-)$ that $s \in \text{vol}(\sigma; \sigma')$.

For the other direction, fix $s \in \text{vol}(\sigma; \sigma')$. We have that, by well-formedness, s can be decomposed as

$$s := s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1}$$

where for each i $s_i \in \sigma; \sigma'$. So there exists s'_i such that $s'_i \upharpoonright_{A,B,-} \in \sigma$ and $s'_i \upharpoonright_{-,B,C} \in \sigma'$.

Let

$$\begin{aligned} t &:= s'_1 \upharpoonright_{A,B,-} \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s'_{n+1} \upharpoonright_{A,B,-} \\ t' &:= s'_1 \upharpoonright_{-,B,C} \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s'_{n+1} \upharpoonright_{-,B,C} \end{aligned}$$

It is immediate from the definition of $\text{vol}(-)$ that $t \in \text{vol}(\sigma)$, $t' \in \text{vol}(\sigma')$, and so it follows that $s \in \text{vol}(\sigma); \text{vol}(\sigma')$.

We also need to show $\text{vol}(\sigma)$ satisfies \downarrow -receptivity. Suppose $s \in \text{vol}(\sigma)$, $\downarrow \in M_{A^\sharp \multimap B^\sharp}^\downarrow$, $s \cdot \downarrow \in P_{A^\sharp \multimap B^\sharp}$, by definition we know s may be decomposed as

$$s = s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1}$$

where for each i , $s_i \in \sigma$ and $r_i \in R$. Since $\epsilon \in \sigma$ so we could set

$$s' := s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1} \cdot \downarrow \cdot \epsilon$$

we obtain by definition that $s' \in \text{vol}(\sigma)$. □

PROPOSITION J.5. *The copycat strategy crashcopy_A is idempotent, i.e.*

$$\text{crashcopy}_A; \text{crashcopy}_A = \text{crashcopy}_A$$

PROOF. Note first that it is immediate from the definition of $\text{crashcopy}_{A^\sharp}$ that

$$\text{crashcopy}_{A^\sharp} = \text{vol}(\text{ccopy}_A)$$

Then, observe that

$$\begin{aligned} \text{crashcopy}_A; \text{crashcopy}_A &= \text{vol}(\text{ccopy}_{A^\sharp}); \text{vol}(\text{ccopy}_{A^\sharp}) && \text{(Def. of crashcopy)} \\ &= \text{vol}(\text{ccopy}_{A^\sharp}; \text{ccopy}_{A^\sharp}) && \text{(Prop. J.4)} \\ &= \text{vol}(\text{ccopy}_{A^\sharp}) && \text{(Prop. ccopy is idempotent)} \\ &= \text{crashcopy}_A && \text{(Def. of crashcopy)} \end{aligned}$$

□

PROPOSITION J.6. *The restriction*

$$\text{vol}(-) : \mathbf{Conc} \rightarrow \mathbf{Crash}$$

of $\text{vol}(-)$ defines a functor.

PROOF. We already argued that

$$\text{crashcopy}_{A^\sharp} = \text{vol}(\text{ccopy}_A)$$

in the proof of Prop. J.5. It remains to show that whenever σ is saturated with respect to ccopy i.e. $\text{ccopy}_A; \sigma; \text{ccopy}_B = \sigma$ then $\text{vol}(\sigma)$ is saturated with respect to crashcopy . For that, just note that:

$$\begin{aligned} \text{crashcopy}_{A^\sharp}; \text{vol}(\sigma); \text{crashcopy}_{B^\sharp} &= \text{vol}(\text{ccopy}_A); \text{vol}(\sigma); \text{vol}(\text{ccopy}_B) \\ &= \text{vol}(\text{ccopy}_A; \sigma; \text{ccopy}_B) \\ &= \text{vol}(\sigma) \end{aligned}$$

□

J.3 Concrete Saturation for Crash-Aware Games

We start by providing the statement of the main result of this section, which concretely characterizes what a saturated strategy in \mathbf{Crash} looks like.

PROPOSITION J.7. *A strategy $\sigma : A \multimap B \in \mathbf{Crash}$ is saturated with respect to crashcopy if and only if it is*

$$\begin{aligned} \text{O-receptive: } \forall s \in \sigma. \forall \alpha \in \Upsilon. \forall m \in M_{A \multimap B}^{\alpha:O}. \exists 1 \leq i \leq \|s\|. \text{epo}_i(s) \cdot m \in P_{A^\Upsilon \multimap B^\Upsilon} \implies \\ \text{epo}_1(s) \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot \text{epo}_i(s) \cdot m \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot \text{epo}_{\|s\|}(s) \in \sigma \end{aligned}$$

$$\rightsquigarrow\text{-closed: } \forall s \in \sigma. \forall t \in P_{A \multimap B}. t \rightsquigarrow_{A \multimap B} s \implies t \in \sigma$$

$$\text{P-delaying: } \forall s \in \sigma. \forall m \in M_{A \multimap B}^P. s = p \cdot m \cdot \frac{1}{2} \cdot t \implies p \cdot \frac{1}{2} \cdot t \in \sigma$$

One might think that it is possible to directly use the proofs of concrete saturation for \mathbf{Conc} from Oliveira Vale et al. [31] across each epoch of a crash-aware strategy to obtain the corresponding concrete saturation result for \mathbf{Crash} . It turns out that those theorems made a key use of prefix-closure of \mathbf{Conc} strategies at specific points that make the proofs not translate to our setting, as strategies in \mathbf{Crash} do not have per-epoch prefix-closure. Because of this, we must reprove some of the results appearing in their appendix, including the Synchronization Lemma.

In the following, we refer to O -receptive, P -delaying closure of a set of plays $S \subseteq P_A$ with $A \in \mathbf{Conc}$ by $\text{dr}(S)$. That is, $\text{dr}(S)$ is the smallest set of plays of P_A such that

$$\forall s \in \text{dr}(S). \forall m \in M_A^O. s \cdot m \in P_A \implies s \cdot m \in \text{dr}(S)$$

$$\forall m \in M_A^P. \forall s \cdot t \in P_A. s \cdot m \cdot t \in \text{dr}(S) \implies s \cdot t \in \text{dr}(S)$$

The following few re-statements of propositions from Oliveira Vale et al. [31] admit essentially the same proofs as might be found there.

PROPOSITION J.8 (SYNCHRONIZATION LEMMA). *Let $s = p \cdot \alpha : m \cdot \alpha' : m' \cdot p'$ be a play of $A \multimap B \in \mathbf{Conc}$. Let $S = \text{dr}(p \cdot m \cdot m' \cdot p')$. Then,*

$$p \cdot m' \cdot m \cdot p' \in \text{ccopy}_A; S; \text{ccopy}_B \iff m' \cdot m \rightsquigarrow_{A \multimap B} m \cdot m'$$

COROLLARY J.9. *Let $s \in P_{A \multimap B}$ with $A \multimap B \in \mathbf{Conc}$ and that t is a play such that*

$$\forall \alpha \in \Upsilon. \pi_\alpha(t) = \pi_\alpha(s)$$

and moreover

$$t \in \text{ccopy}_A; \text{dr}(s); \text{ccopy}_B$$

then,

$$t \rightsquigarrow_{A \rightarrow B} s$$

LEMMA J.10. For every set S of plays of $P_{A \rightarrow B}$:

$$S \subseteq \text{ccopy}_A; S; \text{ccopy}_B$$

The first lemma we need that does require a novel proof is the following.

LEMMA J.11. For any O -receptive and P -delaying set S of plays of $P_{A \rightarrow B}$ with $A \rightarrow B \in \mathbf{Conc}$ it holds that for all $t \in \text{ccopy}_A; S; \text{ccopy}_B$ there exists $s \in S$ such that $t \in \text{ccopy}_A; S; \text{ccopy}_B$ and $\forall \alpha \in \Upsilon. \pi_\alpha(t) = \pi_\alpha(s)$.

PROOF. Fix $t \in \text{ccopy}_A; S; \text{ccopy}_B$ there exists t' such that $t' \upharpoonright_{A,A,-,-} \in \text{ccopy}_A$, $t' \upharpoonright_{-,A,B,-} \in S$ and $t' \upharpoonright_{-, -, B, B} \in \text{ccopy}_B$.

Now, notice that for any $\alpha \in \Upsilon$ there are four possibilities for the lengths of $t' \upharpoonright_{A,A,-,-}$ and $t' \upharpoonright_{-, -, B, B}$.

Both are even-length: It is immediate from the definition of ccopy that $\pi_\alpha(t) = \pi_\alpha(t' \upharpoonright_{-,A,B,-})$. $t' \upharpoonright_{A,A,-,-}$ **is even-length** and $t' \upharpoonright_{-, -, B, B}$ **is odd-length:** Then either $\pi_\alpha(t' \upharpoonright_{-,A,B,-})$ differs from $\pi_\alpha(t)$ by having an extra O -move in the end or by missing a P -move. Either way, we can find a new $t'' \in S$ that either adds the required O -move or removes the missing P -move by O -receptivity or P -delaying respectively.

$t' \upharpoonright_{A,A,-,-}$ **is odd-length** and $t' \upharpoonright_{-, -, B, B}$ **is even-length:** This case is similar to the previous one.

Both length are odd-length: This case is impossible by the switching conditions. □

PROPOSITION J.12. An O -receptive and P -delaying set S of plays $P_{A \rightarrow B}$, for $A \rightarrow B \in \mathbf{Conc}$ is saturated with respect to ccopy if and only if

$$\forall s \in \sigma. \forall t \in P_{A \rightarrow B}. t \rightsquigarrow_{A \rightarrow B} s \implies t \in \sigma$$

PROOF. Suppose S is saturated. It follows that if $s \in S = \text{ccopy}_A; S; \text{ccopy}_B$ and $t \rightsquigarrow_{A \rightarrow B} s$ then there is a sequence of single steps:

$$t = t_0 \rightsquigarrow_{A \rightarrow B} t_1 \rightsquigarrow_{A \rightarrow B} \dots \rightsquigarrow_{A \rightarrow B} t_n = s$$

then by applying the Synchronization Lemma (Prop. J.8) starting with

$$t_{n-1} \rightsquigarrow_{A \rightarrow B} s$$

to conclude that

$$t_{n-1} \in \text{ccopy}_A; \text{dr}(s); \text{ccopy}_B \subseteq \sigma$$

in a finite number of applications we obtain that

$$t = t_0 \in \text{ccopy}_A; \text{dr}(t_1); \text{ccopy}_A \subseteq \text{ccopy}_A; \text{dr}(s); \text{ccopy}_B \subseteq S$$

as desired.

Note that for every set of plays S that satisfies O -receptivity and P -delaying it holds that:

$$S = \bigcup_{s \in S} \text{dr}(s)$$

But

$$\text{ccopy}_A; S; \text{ccopy}_B = \bigcup_{s \in S} \text{ccopy}_A; \text{dr}(s); \text{ccopy}_B$$

by the fact that composition is join-preserving. Hence,

$$t \in \text{ccopy}_A; S; \text{ccopy}_B \iff \exists s \in S. t \in \text{ccopy}_A; \text{dr}(s); \text{ccopy}_B$$

moreover, by J.11, s can be chosen so that

$$\forall \alpha \in \Upsilon. \pi_\alpha(t) = \pi_\alpha(s)$$

by corollary to the Synchronization Lemma (Prop. J.8) it follows that

$$t \in \text{ccopy}_A; \text{dr}(s); \text{ccopy}_B \iff t \rightsquigarrow_{A \rightarrow B} s$$

And hence

$$t \in \text{ccopy}_A; S; \text{ccopy}_B \iff \exists s \in S. t \rightsquigarrow_{A \rightarrow B} s$$

So, suppose $t \in \text{ccopy}_A; S; \text{ccopy}_B$. Then, there is some $s \in \sigma$ such that $t \rightsquigarrow_{A \rightarrow B} s$ and hence by assumption $t \in \sigma$. Hence,

$$\text{ccopy}_A; S; \text{ccopy}_B \subseteq S$$

the reverse containment is exactly lemma J.10 so that it follows that

$$\text{ccopy}_A; S; \text{ccopy}_B = S$$

and hence S is saturated. □

Finally, toward the concrete saturation theorem.

LEMMA J.13. For every strategy $\sigma : A \multimap B \in \underline{\text{Crash}}$ we have

$$\sigma \subseteq \text{crashcopy}_A; \sigma; \text{crashcopy}_B$$

PROOF. Suppose $s \in \sigma$ then we know s can be decomposed as

$$s_1 \cdot \dot{\downarrow} \cdot \dots \cdot \dot{\downarrow} \cdot s_{n+1}$$

where $s_i \in \mathbb{P}_{A^\Upsilon \multimap B^\Upsilon}^{\text{conc}}$. Note then that we have $\text{dr}(s_i) \subseteq \text{ccopy}_{A^\Upsilon}; \text{dr}(s_i); \text{ccopy}_{B^\Upsilon}$. So there exists $t_i \upharpoonright_{A, -, -, B} \in \mathbb{P}_{A^\Upsilon \multimap B^\Upsilon}^{\text{conc}}$ such that $t_i \upharpoonright_{A, A, -, -} \in \text{ccopy}_{A^\Upsilon}$, $t_i \upharpoonright_{-, A, B, -} = s_i$ and $t_i \upharpoonright_{-, -, B, B} \in \text{ccopy}_{B^\Upsilon}$.

So set

$$\begin{aligned} s' &:= t_1 \upharpoonright_{A, A, -, -} \cdot \dot{\downarrow} \cdot \dots \cdot \dot{\downarrow} \cdot t_{n+1} \upharpoonright_{A, A, -, -} \\ s'' &:= t_1 \upharpoonright_{-, A, B, -} \cdot \dot{\downarrow} \cdot \dots \cdot \dot{\downarrow} \cdot t_{n+1} \upharpoonright_{-, A, B, -} \\ s''' &:= t_1 \upharpoonright_{-, -, B, B} \cdot \dot{\downarrow} \cdot \dots \cdot \dot{\downarrow} \cdot t_{n+1} \upharpoonright_{-, -, B, B} \end{aligned}$$

By definition we have $s' \in \text{crashcopy}_A$, $s'' \in \sigma$ and $s''' \in \text{crashcopy}_B$. Hence, $s \in K_{\dot{\downarrow}} \sigma$ □

PROPOSITION J.14. A strategy $\sigma : A \multimap B \in \underline{\text{Crash}}$ is saturated with respect to crashcopy if and only if it is

O-receptive: $\forall s \in \sigma. \forall \alpha \in \Upsilon. \forall m \in M_{A \multimap B}^{\alpha; O}. \exists 1 \leq i \leq \|s\|. \text{epo}_i(s) \cdot m \in P_{A^\Upsilon \multimap B^\Upsilon} \implies$

$$\text{epo}_1(s) \cdot \dot{\downarrow} \cdot \dots \cdot \dot{\downarrow} \cdot \text{epo}_i(s) \cdot m \cdot \dot{\downarrow} \cdot \dots \cdot \dot{\downarrow} \cdot \text{epo}_{\|s\|}(s) \in \sigma$$

\rightsquigarrow -closed: $\forall s \in \sigma. \forall t \in P_{A \multimap B}. t \rightsquigarrow_{A \rightarrow B} s \implies t \in \sigma$

P-delaying: $\forall s \in \sigma. \forall m \in M_{A \multimap B}^P. s = p \cdot m \cdot \dot{\downarrow} \cdot t \implies p \cdot \dot{\downarrow} \cdot t \in \sigma$

PROOF.

(\implies): suppose $\sigma : A \multimap B \in \underline{\text{Crash}}$ is saturated i.e. $\text{crashcopy}_A; \sigma; \text{crashcopy}_B = \sigma$. Let's first show that σ is O-receptive. Fix $s \in \sigma$ by definition s can be decomposed as

$$s = s_1 \cdot \dot{\downarrow} \cdot \dots \cdot \dot{\downarrow} \cdot s_{n+1}$$

Note that for $m \in M_{A \multimap B}^{\alpha; O}$, if $s_i \cdot m = \text{epo}_i(s) \cdot m \in P_{A^\Upsilon \multimap B^\Upsilon}$ then $s_i \cdot m \in \mathbb{P}_{A^\Upsilon \multimap B^\Upsilon}^{\text{conc}}$. Thanks to Oliveira Vale et al. [31] we know there exists u_j for each j such that $u_j \upharpoonright_{A, -, -, B} = s_j \cdot m$

if $j = i$, and otherwise $u_j \upharpoonright_{A,-,-,B} = s_j$, and $u_j \upharpoonright_{A,A,-,-} \in \text{ccopy}_A, u_j \upharpoonright_{-,-,B,B} \in \text{ccopy}_B$ and $u_i \upharpoonright_{-,-,A,B,-} \in \text{dr}(s_i)$.

Let

$$\begin{aligned} u &:= u_1 \upharpoonright_{A,A,-,-} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \upharpoonright_{A,A,-,-} \\ u' &:= u_1 \upharpoonright_{-,-,A,B,-} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \upharpoonright_{-,-,A,B,-} \\ u'' &:= u_1 \upharpoonright_{-,-,B,B} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \upharpoonright_{-,-,B,B} \end{aligned}$$

Then, note that $u \in \text{crashcopy}_A, u' = s, u'' \in \text{crashcopy}_B$ and

$$s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_i \cdot m \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1} \in u; u'; u''$$

so that

$$s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_i \cdot m \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1} \in \text{crashcopy}_A; \sigma; \text{crashcopy}_B = \sigma$$

as σ is saturated.

Now let's show σ is \rightsquigarrow -closed. Fix $s \in \sigma, t \in P_{A \rightarrow B}$ and suppose $t \rightsquigarrow_{A \rightarrow B} s$. More precisely we can decompose s and t as:

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1} \quad t = t_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot t_{n+1}$$

And we have for all $i, t_i \rightsquigarrow_{A \rightarrow B} s_i$. Now we want to show $t \in \sigma$. Since for all i we have $t_i \rightsquigarrow_{A \rightarrow B} s_i$, it follows that $t_i \in \text{ccopy}_A; \text{dr}(s_i); \text{ccopy}_B$ which means there exists u_i such that $u_i \upharpoonright_{A,A,-,-} \in \text{ccopy}_A, u_i \upharpoonright_{-,-,A,B,-} = s_i, u_i \upharpoonright_{-,-,B,B} \in \text{ccopy}_B$ and $u_i \upharpoonright_{A,-,-,B} = t_i$. Let

$$\begin{aligned} u &:= u_1 \upharpoonright_{A,A,-,-} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \upharpoonright_{A,A,-,-} \\ u' &:= u_1 \upharpoonright_{-,-,A,B,-} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \upharpoonright_{-,-,A,B,-} \\ u'' &:= u_1 \upharpoonright_{-,-,B,B} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \upharpoonright_{-,-,B,B} \end{aligned}$$

we have $t \in u; u'; u''$ and $u \in \text{crashcopy}_A, u' \in \text{strat}(s)$ and $u'' \in \text{crashcopy}_B$. So $t \in \text{crashcopy}_A; \text{strat}(s); \text{crashcopy}_B \subseteq \text{crashcopy}_A; \sigma; \text{crashcopy}_B = \sigma$.

Finally, we want to show σ is P -delaying. Fix $s \in \sigma$ such that moreover there is an epoch i and a P -move $m \in M^{\alpha:P}$ such that $s_i = p \cdot m$. By definition s can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

From J.12 we know there exists u_j for each j such that $u_j \upharpoonright_{A,-,-,B} = p$ if $i = j$ otherwise $u_j \upharpoonright_{A,-,-,B} = s_j$, and $u_j \upharpoonright_{A,A,-,-} \in \text{ccopy}_A, u_j \upharpoonright_{-,-,B,B} \in \text{ccopy}_B$ and $u_j \upharpoonright_{-,-,A,B,-} \in \text{dr}(s_j)$

Let

$$\begin{aligned} u &:= u_1 \upharpoonright_{A,A,-,-} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \upharpoonright_{A,A,-,-} \\ u' &:= u_1 \upharpoonright_{-,-,A,B,-} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \upharpoonright_{-,-,A,B,-} \\ u'' &:= u_1 \upharpoonright_{-,-,B,B} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \upharpoonright_{-,-,B,B} \end{aligned}$$

Then, note that $u \in \text{crashcopy}_A, u' = s, u'' \in \text{crashcopy}_B$ and

$$s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot p \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1} \in u; u'; u''$$

so that

$$s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot p \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1} \in K_{\frac{1}{2}} \sigma = \sigma$$

as σ is saturated.

(\Leftarrow): Suppose $\sigma : A \multimap B$ satisfies the O -receptive, P -delaying and \rightsquigarrow -closed conditions. We want to show $\text{crashcopy}_{A; \sigma; \text{crashcopy}_B} = \sigma$.

By lemma J.1 we know that for any $u \in \text{crashcopy}_{A; \sigma; \text{crashcopy}_B}$, there is

$$t = t_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot t_{n+1}$$

such that u can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

where all $s_i \in \text{ccopy}_{A; \text{dr}(t_i); \text{ccopy}_B}$. Since s satisfies the O -receptive, P -delaying and \rightsquigarrow -closed conditions, which means there exists S_i for each i such that $s_i \in S_i$ for all i , S_i satisfy the O -receptive, P -delaying and \rightsquigarrow -closed. Thanks to J.12 we obtain that $\text{ccopy}_{A; \text{dr}(t_i); \text{ccopy}_B} \subseteq S_i$. So we have $s \in \sigma$ so $\text{crashcopy}_{A; \sigma; \text{crashcopy}_B} \subseteq \sigma$

The other direction follows from lemma J.13. □

J.4 Symmetric Monoidal Structure of Crash

PROPOSITION J.15. For any $\sigma : A \multimap B \in \underline{\text{Conc}}$ and $\sigma' : A' \multimap B' \in \underline{\text{Conc}}$ it holds that

$$\text{vol}(\sigma \otimes \sigma') = \text{vol}(\sigma) \otimes \text{vol}(\sigma')$$

PROOF. For one direction, fix $s \in \text{vol}(\sigma) \otimes \text{vol}(\sigma')$. By definition of the tensor s can be decomposed as

$$s := s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

where for each i , $s_i \in \sigma \otimes \sigma'$. It then follows immediately by the definition of $\text{vol}(-)$ that $s \in \text{vol}(\sigma \otimes \sigma')$.

For the other direction, fix $s \in \text{vol}(\sigma \otimes \sigma')$. We have that, by well-formedness, s can be decomposed as

$$s := s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

where for each i , $s_i \in \sigma \otimes \sigma'$. So $s_i \upharpoonright_{A \multimap B} \in \sigma$ and $s_i \upharpoonright_{A' \multimap B'} \in \sigma'$.

Let

$$\begin{aligned} t &:= s_1 \upharpoonright_{A \multimap B} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1} \upharpoonright_{A \multimap B} \\ t' &:= s_1 \upharpoonright_{A' \multimap B'} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1} \upharpoonright_{A' \multimap B'} \end{aligned}$$

It is immediate from the definition of $\text{vol}(-)$ that $t \in \text{vol}(\sigma)$, $t' \in \text{vol}(\sigma')$, and so it follows that $s \in \text{vol}(\sigma) \otimes \text{vol}(\sigma')$. □

LEMMA J.16.

$$\text{crashcopy}_{A \otimes B} = \text{crashcopy}_A \otimes \text{crashcopy}_B$$

PROOF. Note that

$$\begin{aligned} \text{crashcopy}_{A \otimes B} &= \text{vol}(\text{ccopy}_{(A \otimes B)^{\Upsilon}}) && \text{(Def. of crashcopy)} \\ &= \text{vol}(\text{ccopy}_{A^{\Upsilon}} \otimes \text{ccopy}_{B^{\Upsilon}}) && \text{(Symm. Mon. Cat. Conc)} \\ &= \text{vol}(\text{ccopy}_{A^{\Upsilon}}) \otimes \text{vol}(\text{ccopy}_{B^{\Upsilon}}) && \text{(Prop. J.15)} \\ &= \text{crashcopy}_A \otimes \text{crashcopy}_B && \text{(Def. of crashcopy)} \end{aligned}$$

□

LEMMA J.17.

$$- \otimes - : \underline{\text{Crash}} \otimes \underline{\text{Crash}} \rightarrow \underline{\text{Crash}}$$

is a bi-semifunctor

PROOF. For given

$$\begin{aligned} \sigma_1 &: A_1 \multimap A_2, & \sigma_2 &: A_2 \multimap A_3 \\ \tau_1 &: B_1 \multimap B_2, & \tau_2 &: B_2 \multimap B_3 \end{aligned}$$

and fix $s \in (\sigma_1; \sigma_2) \otimes (\tau_1; \tau_2)$ From lemma J.1 we know s can be decomposed as

$$s = s_1 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot s_{n+1}$$

each s_i satisfies

$$\begin{aligned} s_i \upharpoonright_{A_1 \multimap A_3} &\in \text{dr}(t_i); \text{dr}(t'_i) \\ s_i \upharpoonright_{B_1 \multimap B_3} &\in \text{dr}(u_i); \text{dr}(u'_i) \end{aligned}$$

where

$$\begin{aligned} t &:= t_1 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot t_{n+1} \in \sigma_1 \\ t' &:= t_1 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot t'_{n+1} \in \sigma_2 \\ u &:= t_1 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot u_{n+1} \in \tau_1 \\ u' &:= t_1 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot u'_{n+1} \in \tau_2 \end{aligned}$$

futhermore we have for each s_i

$$\begin{aligned} s_i &\in (\text{dr}(t_i); \text{dr}(t'_i)) \otimes (\text{dr}(u_i); \text{dr}(u'_i)) \\ &= (\text{dr}(t_i) \otimes \text{dr}(u_i)); (\text{dr}(t'_i) \otimes \text{dr}(u'_i)) \end{aligned}$$

so there exists s'_i for each i such that

$$\begin{aligned} s'_i \upharpoonright_{A_1 \otimes B_1 \multimap A_3 \otimes B_3} &= s_i \\ s'_i \upharpoonright_{A_1 \otimes B_1 \multimap A_2 \otimes B_2} &\in \text{dr}(t_i) \otimes \text{dr}(u_i) \\ s'_i \upharpoonright_{A_2 \otimes B_2 \multimap A_3 \otimes B_3} &\in \text{dr}(t'_i) \otimes \text{dr}(u'_i) \end{aligned}$$

let

$$\begin{aligned} r &:= s'_1 \upharpoonright_{A_1 \otimes B_1 \multimap A_2 \otimes B_2} \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot s'_{n+1} \upharpoonright_{A_1 \otimes B_1 \multimap A_2 \otimes B_2} \\ r' &:= s'_1 \upharpoonright_{A_2 \otimes B_2 \multimap A_3 \otimes B_3} \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot s'_{n+1} \upharpoonright_{A_2 \otimes B_2 \multimap A_3 \otimes B_3} \end{aligned}$$

we have $r \in \sigma_1 \otimes \tau_1, r' \in \sigma_2 \otimes \tau_2$ futhermore we know $s \in \text{strat}(r); \text{strat}(r') \subseteq (\sigma_1 \otimes \tau_1); (\sigma_2 \otimes \tau_2)$. The other direction is similar.

The enrichment is obvious. First, if $\sigma \subseteq \sigma'$ and $\tau \subseteq \tau'$ it follows immediately from the definition that

$$\sigma \otimes \tau \subseteq \sigma' \otimes \tau'$$

Unions are handled in the same way. □

PROPOSITION J.18. (**Crash**, $- \otimes -, \mathbf{1}$) defines an enriched symmetric monoidal category.

PROOF. We start by showing that the structural morphisms assemble into natural isomorphisms:

$$\begin{array}{ccccc} A \otimes (B \otimes C) & \xrightarrow{\alpha_{A,B,C}} & (A \otimes B) \otimes C & \mathbf{1} \otimes A & \xrightarrow{\lambda_A} & A & A \otimes \mathbf{1} & \xrightarrow{\rho_A} & A \\ \sigma_A \otimes (\sigma_B \otimes \sigma_C) \downarrow & \cong & \downarrow (\sigma_A \otimes \sigma_B) \otimes \sigma_C & \mathbf{1} \otimes \sigma \downarrow & \cong & \downarrow \sigma & \sigma \otimes \mathbf{1} \downarrow & \cong & \downarrow \sigma \\ A' \otimes (B' \otimes C') & \xrightarrow{\alpha_{A',B',C'}} & (A' \otimes B') \otimes C' & \mathbf{1} \otimes B & \xrightarrow{\lambda_B} & B & B \otimes \mathbf{1} & \xrightarrow{\rho_B} & B \end{array}$$

The left and right unital are straight-forward. Indeed, they are given by

$$\lambda''_A := \text{vol}(\lambda_{A^\Upsilon}) \qquad \rho''_A := \text{vol}(\rho_{A^\Upsilon})$$

where λ and ρ are the left and right unitals in **Conc**. It is easy to note that, up to renaming, this makes both unitals the same as crashcopy. Because of this, again, up to renaming, we essentially have:

$$1 \otimes \sigma = \{s' \in P_{(1 \rightarrow 1) \otimes (A \rightarrow B)} \mid \exists s \in \sigma. (\{\epsilon\} \otimes \text{epo}_1(s)) \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot \{(\{\epsilon\} \otimes \text{epo}_{\parallel s \parallel})(s)\} = \sigma$$

which differs from σ only in the shape of the crashes. Therefore, we easily check that:

$$(1 \otimes \sigma); \lambda_B = \sigma; \text{crashcopy}_B = \sigma = \text{crashcopy}_A; \sigma = \lambda_A; \sigma$$

$$(\sigma \otimes 1); \rho_B = \sigma; \text{crashcopy}_B = \sigma = \text{crashcopy}_A; \sigma = \rho_A; \sigma$$

Let β be the braiding in **Conc** and α be the associator in **Conc**. We define the associator and braiding in **Crash** by

$$\beta''_{A,B} := \text{vol}(\beta_{A^\Upsilon, B^\Upsilon}) \quad \text{and} \quad \alpha''_{A,B,C} := \text{vol}(\alpha_{A^\Upsilon, B^\Upsilon, C^\Upsilon})$$

Now, for the associator, the equation essentially follows from the fact that:

$$\begin{aligned} p_\alpha(\sigma_A \otimes (\sigma_B \otimes \sigma_C)); \alpha'_{A', B', C'} &= (p_\alpha(\sigma_A) \otimes (p_\alpha(\sigma_B) \otimes p_\alpha(\sigma_C))); \alpha'_{A', B', C'} \\ &= \alpha'_{A, B, C}; ((p_\alpha(\sigma_A) \otimes p_\alpha(\sigma_B)) \otimes p_\alpha(\sigma_C)) \\ &= \alpha'_{A, B, C}; p_\alpha((\sigma_A \otimes \sigma_B) \otimes \sigma_C) \end{aligned}$$

where for each $s \in \sigma$

$$p_\alpha(s) := \pi'_\alpha(\text{epo}_1(s)) \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot \pi'_\alpha(\text{epo}_{\parallel s \parallel}(s))$$

π'_α is the corresponding projection in **Conc** this is the key step to establish that the naturality square commutes. The reverse direction follows similarly.

For the braiding, let's first show its naturality. For a given strategy $\sigma : A \multimap C, \tau : B \multimap D$ we want to show the following diagram commutes:

$$\begin{array}{ccc} A \otimes B & \xrightarrow{\sigma \otimes \tau} & C \otimes D \\ \beta''_{A,B} \downarrow & & \downarrow \beta''_{C,D} \\ B \otimes A & \xrightarrow{\tau \otimes \sigma} & D \otimes C \end{array}$$

Note that since β is a natural isomorphism, we have that for any $s \in \sigma, t \in \tau$:

$$\begin{array}{ccc} A^\Upsilon \otimes B^\Upsilon & \xrightarrow{\text{dr}(\text{epo}_i(s)) \otimes \text{dr}(\text{epo}_i(t))} & C \otimes D \\ \beta_{A^\Upsilon, B^\Upsilon} \downarrow & & \downarrow \beta_{C^\Upsilon, D^\Upsilon} \\ B^\Upsilon \otimes A^\Upsilon & \xrightarrow{\text{dr}(\text{epo}_i(t)) \otimes \text{dr}(\text{epo}_i(s))} & D^\Upsilon \otimes C^\Upsilon \end{array}$$

commutes.

So for any given $u \in \text{strat}(s) \otimes \text{strat}(t); \beta''_{C,D}$ we have for each i ,

$$\begin{aligned} \text{epo}_i(u) &\in \text{dr}(\text{epo}_i(s)) \otimes \text{dr}(\text{epo}_i(t)); \beta_{C^\Upsilon, D^\Upsilon} \\ &= \beta_{A^\Upsilon, B^\Upsilon}; \text{dr}(\text{epo}_i(t)) \otimes \text{dr}(\text{epo}_i(s)) \end{aligned}$$

And we know for each i , $(\downarrow_A, \downarrow_B) \beta'_{A,B} (\downarrow_B, \downarrow_A)$ and $(\downarrow_C, \downarrow_D) \beta'_{C,D} (\downarrow_D, \downarrow_C)$, so we have

$$s \in \beta'_{A,B}; \tau \otimes \sigma$$

so $\sigma \otimes \tau; \beta''_{C,D} \subseteq \beta''_{A,B}; \tau \otimes \sigma$, other direction is similar. which let us to conclude that

$$\begin{array}{ccc} A \otimes B & \xrightarrow{\sigma \otimes \tau} & C \otimes D \\ \beta''_{A,B} \downarrow & & \downarrow \beta''_{C,D} \\ B \otimes A & \xrightarrow{\tau \otimes \sigma} & D \otimes C \end{array}$$

commutes.

The coherence diagrams follow from functoriality of Vol – together with the fact that Vol – distributes over $-\otimes-$ (Prop. J.15), by noting that all the structural morphisms in Crash were defined by lifting the corresponding structural morphisms in Conc and Rel , which is also why they are isomorphisms. \square

J.5 Crash-Aware Linearizability

PROPOSITION J.19.

$$K_{\downarrow} : \underline{\text{Crash}} \rightarrow \text{Crash}$$

is an enriched oplax semifunctor

PROOF. **Oplax semifunctor:** So we want to show for any $\sigma : A \multimap B$ and $\tau : B \multimap C$ we have

$$K_{\downarrow} (\sigma; \tau) \subseteq K_{\downarrow} \sigma; K_{\downarrow} \tau$$

We have the composition is associative, crashcopy is idempotent and $\sigma \subseteq K_{\downarrow} \sigma$ so

$$\begin{aligned} K_{\downarrow} (\sigma; \tau) &= \text{crashcopy}_A; \sigma; \tau; \text{crashcopy}_C \\ &\subseteq \text{crashcopy}_A; \sigma; \text{crashcopy}_B; \tau; \text{crashcopy}_C \\ &= \text{crashcopy}_A; \sigma; \text{crashcopy}_B; \text{crashcopy}_B; \tau; \text{crashcopy}_C \\ &= K_{\downarrow} \sigma; K_{\downarrow} \tau \end{aligned}$$

Enrichment: Suppose $\sigma \subseteq \sigma'$ then

$$K_{\downarrow} \sigma = \text{crashcopy}; \sigma; \text{crashcopy} \subseteq \text{crashcopy}; \sigma'; \text{crashcopy} = K_{\downarrow} \sigma'$$

by monotonicity of composition. Similarly

$$K_{\downarrow} (\cup_{i \in I} \sigma_i) = \text{crashcopy}; \cup_{i \in I} \sigma_i; \text{crashcopy} = \bigcup_{i \in I} \text{crashcopy}; \sigma_i; \text{crashcopy} = \bigcup_{i \in I} K_{\downarrow} \sigma_i$$

\square

PROPOSITION J.20. For $\tau : A \multimap B \in \underline{\text{Crash}}$,

$$K_{\downarrow} \tau = \{s \in P_{A \rightarrow B} \mid s \text{ is crash-linearizable with respect to } \tau\}$$

PROOF. For one direction, let's fix $s \in K_{\downarrow} \tau$ then by lemma J.1 there exists $t \in \tau$ such that s and t can be decomposed as

$$s = s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1} \quad \text{and} \quad t = t_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot t_{n+1}$$

where for each i we have $s_i \in \text{ccopy}_A; \text{strat}(t_i); \text{ccopy}_B$ which implies that $s_i \sim t_i$ by [31]. Since this is true for all i , we have $s \xrightarrow{\downarrow} \tau$.

For the other direction, fix s crash-linearizable with respect to τ . Then, there exists t in τ such that s and t can be decomposed as

$$s = s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1} \quad \text{and} \quad t = t_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot t_{n+1}$$

and for all i we have $s_i \rightsquigarrow t_i$. So we have $s_i \in \text{ccopy}_A; \text{dr}(t_i); \text{ccopy}_B$ which implies there exists u_i such that $u_i \upharpoonright_{A,A,-} \in \text{ccopy}_A$, $u_i \upharpoonright_{-,A,B,-} \in \text{dr}(t_i)$, $u_i \upharpoonright_{-,-,B,B} \in \text{ccopy}_B$ and $u_i \upharpoonright_{A,-,-} = s_i$. Let

$$\begin{aligned} u &:= u_1 \upharpoonright_{A,A,-} \cdot \downarrow \cdot \dots \cdot \downarrow \cdot u_{n+1} \upharpoonright_{A,A,-} \\ u' &:= u_1 \upharpoonright_{-,A,B,-} \cdot \downarrow \cdot \dots \cdot \downarrow \cdot u_{n+1} \upharpoonright_{-,A,B,-} \\ u'' &:= u_1 \upharpoonright_{-,-,B,B} \cdot \downarrow \cdot \dots \cdot \downarrow \cdot u_{n+1} \upharpoonright_{-,-,B,B} \end{aligned}$$

we have $u \in \text{crashcopy}_A$, $u' \in \text{strat}(t) \subseteq \tau$ and $u'' \in \text{crashcopy}_B$ so we have $s \in K_{\downarrow} \tau$. \square

PROPOSITION J.21. *For $v' : A \in \text{Conc}$ and $v : A \in \underline{\text{Conc}}$, if $v' \rightsquigarrow v$ then $\text{vol}(v') \xrightarrow{\downarrow} \text{vol}(v)$.*

PROOF.

$$\begin{aligned} \text{vol}(v') &= (v' \cdot \downarrow)^* \cdot v' && \text{(Def.)} \\ &\subseteq ((K_{\text{Conc}} v) \cdot \downarrow)^* \cdot K_{\text{Conc}} v && \text{(Lin. equiv. } K_{\text{Conc}}) \\ &= ((v; \text{ccopy}_A) \cdot \downarrow)^* \cdot (v; \text{ccopy}_A) && \text{(Def. of } K_{\text{Conc}}) \\ &= ((v \cdot \downarrow)^* \cdot v); ((\text{ccopy}_A \cdot \downarrow)^* \cdot \text{ccopy}_A) && \text{(Def. } -; - \text{ in } \underline{\text{Crash}}) \\ &= \text{vol}(v); \text{crashcopy}_A = K_{\downarrow} \text{vol}(v) && \text{(Def. vol(-))} \end{aligned}$$

\square

PROPOSITION J.22.

$$K_{\downarrow} \tau = \{s \in P_{A \rightarrow B} \mid s \text{ is crash-linearizable with respect to } \tau\}$$

PROOF. For one direction, let's fix $s \in K_{\downarrow} \tau$ then by lemma J.1 there exists $t \in \tau$ such that s and t can be decomposed as

$$s = s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1} \quad \text{and} \quad t = t_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot t_{n+1}$$

where for each i we have $s_i \in \text{ccopy}_A; \text{strat}(t_i); \text{ccopy}_B$ which implies that $s_i \rightsquigarrow t_i$ by the result from Oliveira Vale et al. [31]. Since this is true for all i , we have $s \xrightarrow{\downarrow} \tau$, as desired.

For the other direction, fix s crash-linearizable with respect to τ . Then, there exists t in τ such that s and t can be decomposed as

$$s = s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1} \quad \text{and} \quad t = t_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot t_{n+1}$$

and for all i we have $s_i \rightsquigarrow t_i$. So we have $s_i \in \text{ccopy}_A; \text{dr}(t_i); \text{ccopy}_B$ which implies there exists u_i such that $u_i \upharpoonright_{A,A,-} \in \text{ccopy}_A$, $u_i \upharpoonright_{-,A,B,-} \in \text{dr}(t_i)$, $u_i \upharpoonright_{-,-,B,B} \in \text{ccopy}_B$ and $u_i \upharpoonright_{A,-,-} = s_i$. Let

$$\begin{aligned} u &:= u_1 \upharpoonright_{A,A,-} \cdot \downarrow \cdot \dots \cdot \downarrow \cdot u_{n+1} \upharpoonright_{A,A,-} \\ u' &:= u_1 \upharpoonright_{-,A,B,-} \cdot \downarrow \cdot \dots \cdot \downarrow \cdot u_{n+1} \upharpoonright_{-,A,B,-} \\ u'' &:= u_1 \upharpoonright_{-,-,B,B} \cdot \downarrow \cdot \dots \cdot \downarrow \cdot u_{n+1} \upharpoonright_{-,-,B,B} \end{aligned}$$

we have $u \in \text{crashcopy}_A$, $u' \in \text{strat}(t) \subseteq \tau$ and $u'' \in \text{crashcopy}_B$ so we have $s \in K_{\downarrow} \tau$. \square

COROLLARY J.23. *$v' : A$ is crash-aware linearizable with respect to $v : A$ if and only if $v' \subseteq K_{\downarrow} v$.*

Now, we move to locality and observational refinement. According to §6 of Oliveira Vale et al. [31], it is enough to prove the following lemma.

LEMMA J.24.

- For any $\sigma : 1 \multimap A \in \mathbf{Crash}$ it holds that $\text{crashcopy}_1; \sigma = \sigma$.
- For $\sigma, \tau : A \multimap B$ and $\sigma', \tau' : A' \multimap B'$ we have

$$\sigma \otimes \sigma' \subseteq \tau \otimes \tau' \implies \sigma \subseteq \sigma' \wedge \tau \subseteq \tau'$$

PROOF.

- For one direction, fix $s \in \text{crashcopy}_1; \sigma$, by J.1 we know s can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

and $s_i \in t_i; u_i$ such that

$$t := t_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot t_{n+1} \in \text{crashcopy}_1$$

$$u := u_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \in \sigma$$

By definition of crashcopy we know $t_i \in \text{ccopy}_1$. Therefore, $t_i = \epsilon$ for all i so $t_i; u_i = u_i$, and therefore $s_i = u_i$. It then follows that $s \in \sigma$.

For the other direction, fix $s \in \sigma$, by definition we know s can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

Now let

$$t := \epsilon \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot \epsilon$$

it is easy to check $t \in \text{crashcopy}_1$. But then, $s \in t; s$, so $s \in \text{crashcopy}_1; \sigma$.

- For given $\sigma, \tau : A \multimap B$ and $\sigma', \tau' : A' \multimap B'$ suppose $\sigma \otimes \sigma' \subseteq \tau \otimes \tau'$. We wish to show that $\sigma \subseteq \tau$. Fix $s \in \sigma$, by well-formedness s can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

By $\frac{1}{2}$ -receptivity it follows that

$$t := \epsilon \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot \epsilon \in \sigma'$$

But $s \otimes t = \{s\}$ up to re-indexing, by definition. So, by monotonicity, $s \in s \otimes t \subseteq \tau \otimes \tau'$. But then it follows that $s \in \tau$, so $\sigma \subseteq \tau$.

The proof of $\sigma' \subseteq \tau'$ is similar.

□

Since by now we have proven all the requirements on the embeddable subcategory of the Karoubi envelope that we have constructed, it follows that

PROPOSITION J.25. *Locality and the equivalence with observational refinement both hold for crash-aware linearizability.*

J.6 Crash Abstraction

We first prove a basic property about $-^b$.

LEMMA J.26. *For $s \in P_{A \multimap B}$,*

- if $\pi_\Upsilon(s \upharpoonright_{-,B}) \in \mathbb{P}_{B^b}$ then $\pi_\Upsilon(s \upharpoonright_{A,-}) \in \mathbb{P}_{A^b}$;
- $\pi_\Upsilon(s) \in \mathbb{P}_{(A \multimap B)^b}$ if and only if $\pi_\Upsilon(s \upharpoonright_{A,-}) \in \mathbb{P}_{A^b}$ and $\pi_\Upsilon(s \upharpoonright_{-,B}) \in \mathbb{P}_{B^b}$.

PROOF.

- Let's prove this result by contradiction. Suppose there exists $s \in P_{A \multimap B}$ such that $\pi_Y(s \upharpoonright_{-,B}) \in \mathbb{P}_{B^b}$ and $\pi_Y(s \upharpoonright_{A,-}) \notin \mathbb{P}_{A^b}$. First of all, by definition s can be decomposed as

$$s = s_1 \cdot \frac{!}{!} \cdot \dots \cdot \frac{!}{!} \cdot s_{n+1}$$

and for each i , $s_i \in \mathbb{P}_{A^Y \multimap B^Y}$

Since $\pi_Y(s \upharpoonright_{A,-}) \notin \mathbb{P}_{A^b}$, there exists a $\alpha \in Y$ such that $\pi_\alpha(s \upharpoonright_{A,-})$ is not a valid sequential play. And, since for each i , $\pi_\alpha(s_i \upharpoonright_A)$ is a valid sequential play, there must exist i such that there is a pending O -move in $\pi_\alpha(s_i \upharpoonright_{A,-})$ and there exists $j > i$ such that $\pi_\alpha(s_j \upharpoonright_{A,-})$ is non-empty. Since $\pi_\alpha(s_i)$ is also a valid sequential play, there must be a pending O -move also in $\pi_\alpha(s_i \upharpoonright_{-,B})$ by the switching condition, and $\pi_\alpha(s_j \upharpoonright_{-,B})$ is also non-empty, since plays are O -starting. But then, this means $\pi_\alpha(s \upharpoonright_{-,B})$ is not a valid sequential play. So $\pi_Y(s \upharpoonright_{-,B}) \notin \mathbb{P}_{B^b}$ which is a contradiction.

- The forward direction is easy to see by definition, while the backward direction follows immediately from the first bullet point. □

COROLLARY J.27.

$$(A \multimap B)^b \cong A^b \multimap B^b$$

We now address the functoriality of $-^b$.

PROPOSITION J.28. $-^b : \mathbf{Crash} \rightarrow \mathbf{Conc}$ defines an enriched oplax semifunctor.

PROOF. Suppose $s \in (\sigma; \tau)^b$. Then, there is $s' \in \text{int}(\sigma, \tau)$ such that $s = \pi_Y(s' \upharpoonright_{A,-,C})$ and in particular $\pi_Y(s' \upharpoonright_{A,-,C}) \in \mathbb{P}_{A \multimap B^b}$. Note then that by Prop. J.26 it follows that $\pi_Y(s' \upharpoonright_{-,B,C} \upharpoonright_{-,C}) = \pi_Y(s' \upharpoonright_{A,-,C} \upharpoonright_{-,C}) \in \mathbb{P}_{C^b}$, so that by the same proposition it follows that $\pi_Y(s' \upharpoonright_{-,B,C} \upharpoonright_{B,-}) \in P_{B^b}$ and therefore $\pi_Y(s' \upharpoonright_{-,B,C}) \in \mathbb{P}_{(B \multimap C)^b}$. At this point we have $s' \upharpoonright_{A,B,-} \upharpoonright_{-,B} = s' \upharpoonright_{-,B,C} \upharpoonright_{B,-} \in \mathbb{P}_{B^b}$ so that analogous reasoning gives that $s' \upharpoonright_{A,B,-} \in \mathbb{P}_{(A \multimap B)^b}$. But we also have $s' \upharpoonright_{A,B,-} \in \sigma$ and $s' \upharpoonright_{-,B,C} \in \tau$, so that we have shown that $\pi_Y(s' \upharpoonright_{A,B,-}) \in \sigma^b$ and $\pi_Y(s' \upharpoonright_{-,B,C}) \in \tau^b$. We then obtain that $\pi_Y(s') \upharpoonright_{A^b, B^b, -} = \pi_Y(s' \upharpoonright_{A^b, B^b, -}) \in \sigma^b$ and $\pi_Y(s') \upharpoonright_{-, B^b, C^b} = \pi_Y(s' \upharpoonright_{-, B, C}) \in \tau^b$. Hence, $\pi_Y(s') \in \text{int}(\sigma^b, \tau^b)$ and therefore $s = \pi_Y(s' \upharpoonright_{A,-,C}) = \pi_Y(s') \upharpoonright_{A^b, -, C^b} \in \sigma^b; \tau^b$.

We move on to the enrichment. Suppose $\sigma \subseteq \sigma'$, fix $s \in \sigma^b$, by definition there exists $s' \in \sigma$ such that $\pi_Y(s') = s$. Since $\sigma \subseteq \sigma'$ so $s' \in \sigma'$ so $s \in (\sigma')^b$. So we have $\sigma^b \subseteq (\sigma')^b$

Given a family of stratgies $(\sigma_i : A \multimap B)_{i \in I}$. For one direction, fix $s \in (\cup_{i \in I} \sigma_i)^b$ so there exists $s' \in \cup_{i \in I} \sigma_i$ such that $\pi_Y(s') = s$. So there exists $i \in I$ such that $s' \in \sigma_i$ and $\pi_Y(s') = s \in \mathbb{P}_{A^b}$ which means $s \in \sigma_i^b$. So $s \in \cup_{i \in I} \sigma_i^b$. For the other direction, fix $s \in \cup_{i \in I} \sigma_i^b$, we know there exists $i \in I$ such that $s' \in \sigma_i$ and $s' \upharpoonright_{-,B} = s$. So $s' \in \cup_{i \in I} \sigma_i$ and $\pi_Y(s') = s \in \mathbb{P}_{A^b}$, so $s \in (\cup_{i \in I} \sigma_i)^b$ □

PROPOSITION J.29. $(\text{crashcopy}_A)^b = \text{ccopy}_{A^b}$

PROOF. By definition it is easy to see $\text{ccopy}_{A^b} \subseteq (\text{crashcopy}_A)^b$.

For the other direction, fix $s \in (\text{crashcopy}_A)^b$, by definition, s can be decomposed as

$$s = s_1 \cdot s_2 \cdot \dots \cdot s_{n+1}$$

where for each i , $s_i \in \text{ccopy}_{A^b}$.

By the definition of copy we know for any two plays $t, t' \in \text{copy}$, if $t \cdot t'$ is a valid play, then $t \cdot t' \in \text{copy}$. So for any $\alpha \in Y$ we get that $\pi_\alpha(s)$ is in copy_{A^b} so $s \in \text{ccopy}_{A^b}$ □

COROLLARY J.30. For every $\sigma : A \multimap B \in \mathbf{Crash}$:

$$(K_{\frac{!}{!}} \sigma)^b \subseteq K_{\text{Conc}} \sigma^b$$

We now move to the functoriality like properties of re-crash operation $-^\#$.

PROPOSITION J.31. *For strategies $\sigma : A^b \multimap B^b$, and $\tau : B^b \multimap C^b$, the following hold:*

- if $\sigma \subseteq \sigma'$, for $\sigma' : A^b \multimap B^b$, then $\sigma^\# \subseteq (\sigma')^\#$
- Given a family $(\sigma_i : A^b \multimap B^b)_{i \in I}$ it holds that $(\cup_{i \in I} \sigma_i)^\# = \cup_{i \in I} \sigma_i^\#$
- $\sigma^\#; \tau^\# \subseteq (\sigma; \tau)^\#$

PROOF.

- Suppose $\sigma \subseteq \sigma'$, fix $s \in \sigma^\#$ by definition we know $\pi_Y(s) \in \sigma$. Since $\sigma \subseteq \sigma'$, so $\pi_Y(s) \in \sigma'$ which implies $s \in (\sigma')^\#$. So $\sigma^\# \subseteq (\sigma')^\#$
- For a given family $(\sigma_i)_{i \in I}$. For one direction, fix $s \in (\cup_{i \in I} \sigma_i)^\#$ we know $\pi_Y(s) \in \cup_{i \in I} \sigma_i$ which means there exists $i \in I$ such that $\pi_Y(s) \in \sigma_i$. By definition this means $s \in \sigma_i^\#$, so $s \in \cup_{i \in I} \sigma_i^\#$. For the other direction, let fix $s \in \cup_{i \in I} \sigma_i^\#$, by definition, there exists $i \in I$ such that $\pi_Y(s) \in \sigma_i \subseteq \cup_{i \in I} \sigma_i$. So $s \in \cup_{i \in I} \sigma_i^\#$
- fix $s \in \sigma^\#; \tau^\#$, by definition there exists $s' \in \text{int}(\sigma^\#, \tau^\#)$ such that $s' \upharpoonright_{A,-,C} = s$, $s' \upharpoonright_{A,B,-} \in \sigma^\#$ and $s' \upharpoonright_{-,B,C} \in \tau^\#$.

Set $t := \pi_Y(s' \upharpoonright_{A,B,-})$, $u := \pi_Y(s' \upharpoonright_{-,B,C})$, it is easy to check

$$\pi_Y(s) = \pi_Y(s' \upharpoonright_{A,-,C}) \in \text{strat}(t); \text{strat}(u) \subseteq \sigma; \tau$$

so $s \in (\sigma; \tau)^\#$. Well-formedness of $\pi_Y(s)$ is guaranteed by well-formedness of $\pi_Y(s' \upharpoonright_{A,-,-})$, $\pi_Y(s' \upharpoonright_{-, -, C})$ by Prop. J.26. □

PROPOSITION J.32. *For all $A \in \text{Crash}$,*

$$\text{ccopy}_{A^b}^\# \subseteq \text{crashcopy}_A$$

PROOF. Fix $s \in \text{ccopy}_{A^b}^\#$, by definition we know there exists $s' \in \text{ccopy}_{A^b}$ such that s can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

where $s_1 \cdot \dots \cdot s_{n+1} = s'$ and for each i , s_i is a play in \mathbb{P}_{A^Y} . Furthermore we know for each $i \neq n+1$, $\alpha \in Y$, there is no pending O moves in $\pi_\alpha(s_i)$.

Now we want to show for all i , $\alpha \in Y$, $p \sqsubseteq_{\text{even}} \pi_\alpha(s_i)$ we have $p \upharpoonright_{A,-} = p \upharpoonright_{-,A}$ by contradiction. Suppose i is the smallest i such that there exists $\alpha \in Y$, $p \sqsubseteq_{\text{even}} \pi_\alpha(s_i)$ such that $p \upharpoonright_{A,-} \neq p \upharpoonright_{-,A}$.

Since $s'' := s_1 \cdot \dots \cdot s_i \sqsubseteq s'$ so we know $s'' \in \text{ccopy}_{A^b}$ which means for every $q \sqsubseteq_{\text{even}} \pi_\alpha(s'')$ we have $q \upharpoonright_{A,-} = q \upharpoonright_{-,A}$. Since there is no pending O moves in s_1, \dots, s_{i-1} so $p' := \pi_\alpha(s_1 \cdot \dots \cdot s_{i-1})$ has to be an even prefix of $\pi_\alpha(s'')$ so is $p' \cdot p$. But it is easy to see $p' \cdot p \upharpoonright_{A,-} \neq p' \cdot p \upharpoonright_{-,A}$ which is a contradiction. So we know for each i , $s_i \in \text{crashcopy}_{A^b}$. By definition, $s \in \text{crashcopy}_A$. □

We also take the opportunity to show that how $-^\#$ interacts with horizontal composition.

LEMMA J.33. *For any $\sigma : (A_1)^b \multimap (B_1)^b$, $\tau : (A_2)^b \multimap (B_2)^b$,*

$$(\sigma \otimes \tau)^\# = \sigma^\# \otimes \tau^\#$$

PROOF. For one direction fix $s \in (\sigma \otimes \tau)^\#$, we know s can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

for each i we know there exists t_i, u_i such that $s_i \in t_i \otimes u_i$ such that $u_1 \cdot \dots \cdot u_{n+1} \in \sigma$ and $t_1 \cdot \dots \cdot t_{n+1} \in \tau$

By definition we can see

$$\begin{aligned} t &:= t_1 \cdot \underline{\downarrow} \cdot \dots \cdot \underline{\downarrow} \cdot t_{n+1} \in \sigma^\sharp \\ u &:= u_1 \cdot \underline{\downarrow} \cdot \dots \cdot \underline{\downarrow} \cdot u_{n+1} \in \tau^\sharp \end{aligned}$$

so we know $s \in \text{strat}(t) \otimes \text{strat}(u) \subseteq \sigma^\sharp \otimes \tau^\sharp$

For the other direction let's fix $s \in \sigma^\sharp \otimes \tau^\sharp$ then there exists $t \in \sigma^\sharp, u \in \tau^\sharp$ such that $s \in \text{strat}(t) \otimes \text{strat}(u)$. Let $s' = \pi_\Upsilon(t) \otimes \pi_\Upsilon(u)$ by definition we know $s' \in \sigma \otimes \tau$ and $\pi_\Upsilon(s) = s' \otimes s \in (\sigma \otimes \tau)^\sharp$ \square

J.7 Strict Linearizability

We start by showing some auxiliary lemmas.

LEMMA J.34. *For any given $\sigma, \tau \in \text{Conc}$ we have*

$$\sigma^\sharp; \text{crashcopy}_B; \tau^\sharp = \sigma^\sharp; \text{ccopy}_{B^b}^\sharp; \tau^\sharp$$

PROOF. One direction follow from an already proved lemma.

$$\sigma^\sharp; \text{ccopy}_{B^b}^\sharp; \tau^\sharp \subseteq \sigma^\sharp; \text{crashcopy}_B; \tau^\sharp$$

Now let's try to prove the other direction by contradiction. Suppose $s \in \sigma^\sharp; \text{crashcopy}_B; \tau^\sharp$, assume $s \notin \sigma^\sharp; \text{ccopy}_{B^b}^\sharp; \tau^\sharp$. So there exists s' such that $s' \upharpoonright_{A,B,-} \in \sigma^\sharp, s' \upharpoonright_{-,B,B,-} \in \text{crashcopy}_B$ but $s' \upharpoonright_{-,B,B,-} \notin \text{ccopy}_{B^b}^\sharp, s' \upharpoonright_{-, -, B, C} \in \tau^\sharp$ and $s' \upharpoonright_{A,-, -} = s$

Since we already know $(\text{crashcopy})^b = \text{ccopy}$, it follows that $\text{ccopy}_{B^b}^\sharp = (\text{crashcopy}^b)^\sharp$. So we get that $s' \upharpoonright_{-,B,B,-} \notin (\text{crashcopy}^b)^\sharp$. But this means that at least one of $\pi_\Upsilon(s' \upharpoonright_{-,B,-,-})$ and $\pi_\Upsilon(s' \upharpoonright_{-, -, B, -})$ is not well-formed. But we know both $s' \upharpoonright_{A,B,-}$ and $s' \upharpoonright_{-, -, B, C}$ are well-formed which is a contradiction. \square

LEMMA J.35. *For any given $\sigma, \tau \in \text{Conc}$ we have*

$$\sigma^\sharp; \text{crashcopy}_B; \tau^\sharp = \sigma^\sharp; \tau^\sharp$$

PROOF. For one direction, note that $\sigma^\sharp; \tau^\sharp \subseteq \sigma^\sharp; \text{crashcopy}_B; \tau^\sharp$ follows from what we showed about K_{\downarrow}^- .

For the other direction notice that

$$\begin{aligned} \sigma^\sharp; \text{crashcopy}_B; \tau^\sharp &= \sigma^\sharp; \text{ccopy}_{B^b}^\sharp; \tau^\sharp \\ &\subseteq \sigma^\sharp; (\text{ccopy}_{B^b}; \tau)^\sharp \\ &= \sigma^\sharp; \tau^\sharp \end{aligned}$$

\square

PROPOSITION J.36.

$$\text{crashcopy}_A; \text{ccopy}_{A^b}^\sharp = \text{ccopy}_{A^b}^\sharp$$

PROOF. One direction $\text{ccopy}_{A^b}^\sharp \subseteq \text{crashcopy}_A; \text{ccopy}_{A^b}^\sharp$ follows readily from what we showed about K_{\downarrow}^- .

The other direction, let's fix $s \in \text{crashcopy}_A; \text{ccopy}_{A^b}^\sharp$, then there exists s' such that $s' \upharpoonright_{A,A,-} \in \text{crashcopy}_A, s' \upharpoonright_{-,A,A} \in \text{ccopy}_{A^b}^\sharp$, and $s' \upharpoonright_{A,-,A} = s$.

Let's first show that $\pi_\alpha(s)$ is well-formed by contradiction. By definition we know s can be decomposed as

$$s = s_1 \cdot \downarrow \cdot \dots \cdot \downarrow \cdot s_{n+1}$$

If s is not well-formed, we know there exists $\alpha \in \Upsilon$ such that there exists l and $m > l$ such that $\pi_\alpha(s_l \cdot s_m)$ is not well-formed and $\pi_\alpha(s_m) \neq \epsilon$. Let m be the first m satisfies the property.

This means either $\pi_\alpha(s_l \cdot s_m) \uparrow_{A,-}$ or $\pi_\alpha(s_l \cdot s_m) \uparrow_{-,A}$ is not well-formed. But we know $\pi_\alpha(s_l \cdot s_m) \uparrow_{-,A} = \pi_\alpha(s'_l \cdot s'_m) \uparrow_{-,A}$ which is well-formed by definition. So the only possibility left is $\pi_\alpha(s_l \cdot s_m) \uparrow_{A,-}$ is not well-formed.

This means there is a pending O move in $\pi_\alpha(s_l) \uparrow_{A,-}$ which will become a P move in s_l and this will force $\pi_\alpha(s_l) \uparrow_{-,A}$ has a pending O move. But we already know $\pi_\alpha(s_m) \uparrow_{-,A}$ is non-empty, so this is a contradiction.

Now we want to show $\pi_\Upsilon(s) \in \text{ccopy}_{A^b}$. We only need to show for each $\alpha \in \Upsilon$ we have $\pi_\alpha(s) \in \text{copy}_{(A^b)^\alpha}$. Let's also show this by contradiction, suppose there exists $\alpha \in \Upsilon$ and $p \sqsubseteq_{\text{even}} \pi_\alpha(s)$ such that $\pi_\alpha(p) \uparrow_{A,-} \neq \pi_\alpha(p) \uparrow_{-,A}$

So there exists $p' \sqsubseteq \pi_\alpha(s')$ such that $p' \uparrow_{A,-,A} = p$. Since $p' \uparrow_{A,A,-}$ is in $\text{copy}_{(A^b)^\alpha}$ so $p' \uparrow_{A,-,-} = p' \uparrow_{-,A,-}$. Same we can get $p' \uparrow_{-,A,-} = p' \uparrow_{-, -,A}$. So we know $p' \uparrow_{A,-,-} = p' \uparrow_{-, -,A}$ which means $p \uparrow_{A,-} = p \uparrow_{-,A}$ which is a contradiction. \square

With all these lemmas proved, we are ready to show observational refinement.

PROPOSITION J.37. *Let $\sigma : A^b \multimap B^b \in \mathbf{Conc}$ and $\tau : B^b \multimap C^b \in \mathbf{Conc}$, then*

$$\text{str}(\sigma); \text{str}(\tau) \subseteq \text{str}(\sigma; \tau)$$

PROOF.

$$\begin{aligned} \text{str}(\sigma); \text{str}(\tau) &= K_{\frac{1}{2}} \sigma^\#; K_{\frac{1}{2}} \tau^\# \\ &= \text{crashcopy}_A; \sigma^\#; \text{crashcopy}_B; \tau^\#; \text{crashcopy}_C \\ &= \text{crashcopy}_A; \sigma^\#; \tau^\#; \text{crashcopy}_C \\ &\subseteq \text{crashcopy}_A; (\sigma; \tau)^\#; \text{crashcopy}_C \\ &\subseteq K_{\frac{1}{2}} (\sigma; \tau)^\# \\ &= \text{str}(\sigma; \tau) \end{aligned}$$

\square

PROPOSITION J.38. *If $v'_A \subseteq \text{str}(v_A)$ then, for all $\sigma : A^b \multimap B^b \in \mathbf{Conc}$ that implements an object linearizable to $v_B : B^b$ using v_A , i.e.*

$$v_A; \sigma \subseteq v_B$$

It holds that,

$$v'_A; \text{str}(\sigma) \subseteq \text{str}(v_B)$$

PROOF. For the forward direction, we start by noting that since

$$v_A; \sigma \subseteq v_B$$

First, note that

$$v'_A \subseteq \text{str}(v_A) \subseteq \text{str}(K_{\text{Conc}} v_A)$$

By the observational refinement property on \mathbf{Conc} it follows that

$$K_{\text{Conc}} v_A; \sigma \subseteq v_B$$

Then, by Prop. J.37

$$v'_A; \text{str}(\sigma) \subseteq \text{str}(K_{\text{Conc}} v_A); \text{str}(\sigma) = \text{str}(K_{\text{Conc}} v_A; \sigma) \subseteq \text{str}(v_B)$$

□

For locality, note first that:

PROPOSITION J.39. For $\sigma : A_1^b \multimap B_1^b$ and $\tau : A_2^b \multimap B_2^b$,

$$\text{str}(\sigma \otimes \tau) = \text{str}(\sigma) \otimes \text{str}(\tau)$$

PROOF.

$$\text{str}(\sigma \otimes \tau) = K_{\sharp} (\sigma \otimes \tau)^{\sharp} = K_{\sharp} (\sigma^{\sharp} \otimes \tau^{\sharp}) = K_{\sharp} \sigma^{\sharp} \otimes K_{\sharp} \tau^{\sharp} = \text{str}(\sigma) \otimes \text{str}(\tau)$$

□

PROPOSITION J.40 (LOCALITY). For $v'_A : A, v'_B : B \in \mathbf{Crash}$ and $v_A : A, v_B : B \in \mathbf{Conc}$:

$$v'_A \subseteq \text{str}(v_A) \text{ and } v'_B \subseteq \text{str}(v_B) \text{ if and only if } v'_A \otimes v'_B \subseteq \text{str}(v_A \otimes v_B)$$

PROOF. For the forward direction we have that by monotonicity:

$$v'_A \otimes v'_B \subseteq \text{str}(v_A) \otimes \text{str}(v_B) = \text{str}(v_A \otimes v_B)$$

For the reverse direction, first note that

$$v'_A \otimes v'_B \subseteq \text{str}(v_A \otimes v_B) = \text{str}(v_A) \otimes \text{str}(v_B)$$

since we have shown the tensor is an order-isomorphism the result follows. □

J.8 Durability

LEMMA J.41. Let $s \in P_{A \multimap B}$ for $A, B \in \mathbf{Crash}$.

- If $s \upharpoonright_{-,B}$ is durable then $s \upharpoonright_{A,-}$ is durable.
- s is durable if and only if $s \upharpoonright_{A,-}$ and $s \upharpoonright_{-,B}$ are both durable.

PROOF.

- Let's prove this by contradiction. Suppose $s \upharpoonright_{-,B}$ is durable, but $s \upharpoonright_{A,-}$ is not. Then there exists $i \neq j$ such that $\exists \alpha \in Y. \alpha \in Y(\text{epo}_i(s \upharpoonright_{A,-})) \cap Y(\text{epo}_j(s \upharpoonright_{A,-}))$. Since both $\text{epo}_i(s), \text{epo}_j(s) \in \mathbb{P}_{A^Y \multimap B^Y}$, it follows that $\pi_\alpha(\text{epo}_i(s)), \pi_\alpha(\text{epo}_j(s)) \in \mathbb{P}_{A^\alpha \multimap B^\alpha}$. So by the switching condition, and since $\pi_\alpha(\text{epo}_i(s) \upharpoonright_{A,-})$ is non-empty we have that $\pi_\alpha(\text{epo}_i(s) \upharpoonright_{-,B})$ is also non-empty. The same applies for j . Again by the switching condition, $\alpha \in Y(\text{epo}_i(s \upharpoonright_{-,B})) \cap Y(\text{epo}_j(s \upharpoonright_{-,B}))$. But we know $s \upharpoonright_{-,B}$ is durable which is a contradiction.
- (\Rightarrow) It is easy to see by definition
- (\Leftarrow) Suppose both $s \upharpoonright_{A,-}$ and $s \upharpoonright_{-,B}$ are durable. Let's prove s is also durable by contradiction. Suppose s is not, so there exists $i \neq j$ such that $\exists \alpha \in Y. \alpha \in Y(\text{epo}_i(s)) \cap Y(\text{epo}_j(s))$. So $\pi_\alpha(\text{epo}_i(s) \upharpoonright_{-,B})$ and $\pi_\alpha(\text{epo}_j(s) \upharpoonright_{-,B})$ both can't be empty play. So $\alpha \in Y(\text{epo}_i(s) \upharpoonright_{-,B}) \cap Y(\text{epo}_j(s) \upharpoonright_{-,B})$ by the switching condition. But we know $s \upharpoonright_{-,B}$ is durable which is a contradiction.

□

PROPOSITION J.42. Durable strategies compose.

PROOF. Suppose $\sigma : A \multimap B, \tau : B \multimap C$, are both durable.

We already have that $\sigma; \tau : A \multimap C$ is well-defined, it remains to show that it is also durable. By definition of composition we know for any $s \in \sigma; \tau$ there exists $t \in \sigma, u \in \tau$ such that $s \upharpoonright_{A,-} = t \upharpoonright_{A,-}$ and $s \upharpoonright_{-,C} = u \upharpoonright_{-,C}$.

Since both σ, τ are durable, by proposition J.41 $t \upharpoonright_{A,-}$ and $u \upharpoonright_{-,C}$ are both durable. So thanks to proposition J.41 again we know s is also durable. \square

This means that the restriction of **Crash** to durable strategies, **Dur**, defines a semicategory.

LEMMA J.43. *For sets of plays*

$$S \subseteq P_{A \multimap B} \quad T \subseteq P_{B \multimap C}$$

and

$$(S \cap P_{A \multimap B}^{\text{dur}}); (T \cap P_{B \multimap C}^{\text{dur}}) = (S; T) \cap P_{A \multimap C}^{\text{dur}}$$

PROOF. For one direction, fix $s \in (S \cap P_{A \multimap B}^{\text{dur}}); (T \cap P_{B \multimap C}^{\text{dur}})$, we know there exists s' such that $s' \upharpoonright_{A,B,-} \in S \cap P_{A \multimap B}^{\text{dur}}, s' \upharpoonright_{-,B,C} \in T \cap P_{B \multimap C}^{\text{dur}}$ and $s' \upharpoonright_{A,-,C} = s$.

In particular we know $s' \upharpoonright_{A,B,-} \in S$ and $s' \upharpoonright_{-,B,C} \in T$ which implies $s \in S; T$. By applying proposition J.41 we know $s \upharpoonright_{A,-} = s' \upharpoonright_{A,-,-}, s \upharpoonright_{-,C} = s' \upharpoonright_{-,,-,C}$ are both durable. So also by proposition J.41 we know s is durable. So $s \in (S; T) \cap P_{A \multimap C}^{\text{dur}}$.

For other direction, fix $s \in (S; T) \cap P_{A \multimap C}^{\text{dur}}$. We know there exists s' such that $s' \upharpoonright_{A,B,-} \in S, s' \upharpoonright_{-,B,C} \in T, s' \upharpoonright_{A,-,C} = s$. By proposition J.41 we $s \upharpoonright_{-,C}$ is durable so is $s' \upharpoonright_{-,,-,C}$. By proposition J.41 we know $s' \upharpoonright_{-,B,-}$ has to be durable. Then by applying proposition J.41 again we get $s' \upharpoonright_{A,-,-}$ is durable.

So we know $s' \upharpoonright_{A,B,-} \in S \cap P_{A \multimap B}^{\text{dur}}$ and $s' \upharpoonright_{-,B,C} \in T \cap P_{B \multimap C}^{\text{dur}}$. \square

COROLLARY J.44. *The assignment:*

$$A \longmapsto A \quad \sigma : A \multimap B \longmapsto \sigma \cap P_{A \multimap B}^{\text{dur}}$$

defines a semifunctor from **Crash** to **Dur**.

COROLLARY J.45. *durcopy is idempotent.*

COROLLARY J.46. *For any given (lax, oplax, semi) functor $F : \mathbf{C} \rightarrow \mathbf{Crash}$, we have*

$$F(-) \cap P_{F X \multimap F Y}^{\text{dur}} : \mathbf{C} \rightarrow \mathbf{Dur}$$

also defines (lax, oplax, semi) functor respectively

PROPOSITION J.47. *A strategy $\sigma : A \multimap B \in \mathbf{Crash}$ is saturated with respect to durcopy if and only if it is a durable strategy and*

durably O-receptive:

$$\forall s \in \sigma. \forall \alpha \in \Upsilon. \forall m \in M_{A \multimap B}^{\alpha; O}. \exists i \leq \|s\|.$$

$$\text{epo}_i(s) \cdot m \in P_{A^\Upsilon \multimap B^\Upsilon} \wedge \forall j \neq i. \Upsilon(m) \notin \Upsilon(\text{epo}_j(s)) \implies$$

$$\text{epo}_1(s) \cdot \frac{1}{2} \dots \frac{1}{2} \cdot \text{epo}_i(s) \cdot m \cdot \frac{1}{2} \dots \frac{1}{2} \cdot \text{epo}_{\|s\|}(s) \in \sigma$$

\rightsquigarrow -closed: $\forall s \in \sigma. \forall t \in P_{A \multimap B}. t \rightsquigarrow_{A \multimap B} s \implies t \in \sigma$

P -delaying: $\forall s \in \sigma. \forall m \in M_{A \multimap B}^P. \forall m_\frac{1}{2} \in M_{A \multimap B}^{\frac{1}{2}}. s = p \cdot m \cdot m_\frac{1}{2} \cdot t \implies p \cdot m_\frac{1}{2} \cdot t \in \sigma$

PROOF. For one direction, notice that

$$\text{durcopy}_A; \sigma; \text{durcopy}_B = (\text{crashcopy}_A; \sigma; \text{crashcopy}_B) \cap P_{A \multimap B}^{\text{dur}}$$

By J.7 we know $\text{crashcopy}_A; \sigma; \text{crashcopy}_B$ satisfies O -receptive, P -delaying and \rightsquigarrow -closed. So $(\text{crashcopy}_A; \sigma; \text{crashcopy}_B) \cap P_{A \multimap B}^{\text{dur}}$ will satisfy durably O -receptive and \rightsquigarrow -closed.

For the other direction, suppose σ is durable, durably O -receptive, P -delaying and \rightsquigarrow -closed. We want to show $\sigma = \text{durcopy}_A; \sigma; \text{durcopy}_B$.

Set σ' be the smallest strategy contain σ and satisfies O -receptive and P -delaying. By definition we know $\sigma' \cap P_{A \multimap B}^{\text{dur}} = \sigma$ and σ' satisfies \rightsquigarrow -closed.

Since σ' satisfies O -receptive, P -delaying and \rightsquigarrow -closed, $\sigma' = \text{crashcopy}_A; \sigma'; \text{crashcopy}_B$. So we have $(\text{crashcopy}_A; \sigma'; \text{crashcopy}_B) \cap P_{A \multimap B}^{\text{dur}} = \sigma$.

Now we have

$$\begin{aligned} \sigma &= \sigma' \cap P_{A \multimap B}^{\text{dur}} \\ &= (\text{crashcopy}_A; \sigma'; \text{crashcopy}_B) \cap P_{A \multimap B}^{\text{dur}} \\ &= (\text{crashcopy}_A \cap P_{A \multimap A}^{\text{dur}}); (\sigma' \cap P_{A \multimap B}^{\text{dur}}); (\text{crashcopy}_B \cap P_{B \multimap B}^{\text{dur}}) \\ &= \text{durcopy}_A; \sigma; \text{durcopy}_B \end{aligned}$$

□

PROPOSITION J.48. For any $A, B \in \underline{\text{Crash}}$ and $\sigma : A^b \multimap B^b \in \underline{\text{Conc}}$, $\text{dur}(\sigma) \in \underline{\text{Dur}}$.

PROOF. We want to show $\text{dur}(\sigma) = \text{durcopy}_A; K_{\text{Conc}} \sigma^\#; \text{durcopy}_B$. For the direction $\text{dur}(\sigma) \subseteq \text{durcopy}_A; \text{dur}(\sigma); \text{durcopy}_B$ follows from the analogous fact about $K_{\text{Conc}} -$ and monotonicity of $- \cap P_{A \multimap B}^{\text{dur}}$.

Now for the other direction, fix $s \in \text{durcopy}_A; \text{dur}(\sigma); \text{durcopy}_B$ first by definition s can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

and for each i we know $s_i \in \text{ccopy}_A; s_i; \text{ccopy}_B$ and we know for each $i \neq j$ we have $\Upsilon(s_i) \neq \Upsilon(s_j)$

So there exists s'_i such that $s'_i \upharpoonright_{A,A,-,-} \in \text{ccopy}_A$, $s'_i \upharpoonright_{-,A,B,-} = s_i$, $s'_i \upharpoonright_{-,-,B,B} \in \text{ccopy}_B$. Since we already know for each $i \neq j$ we have $\Upsilon(s_i) \neq \Upsilon(s_j)$ so we know $\Upsilon(s'_i) \neq \Upsilon(s'_j)$.

So we know

$$\begin{aligned} t &:= s'_1 \upharpoonright_{A,A,-,-} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s'_{n+1} \upharpoonright_{A,A,-,-} \\ t' &:= s'_1 \upharpoonright_{-,A,B,-} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s'_{n+1} \upharpoonright_{-,A,B,-} \\ t'' &:= s'_1 \upharpoonright_{-,-,B,B} \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s'_{n+1} \upharpoonright_{-,-,B,B} \end{aligned}$$

such that $\pi_\Upsilon(t) \in \text{ccopy}_{A^b}$, $\pi_\Upsilon(t'') \in \text{ccopy}_{B^b}$ and $\pi_\Upsilon(t') \in \sigma$, so $s \in \text{dur}(\sigma)$.

For the $\frac{1}{2}$ -receptivity, notice that suppose $s \in \text{dur}(\sigma)$, $m_{\frac{1}{2}} \in M_{A \multimap B}^{\frac{1}{2}}$, $s \cdot m_{\frac{1}{2}} \in P_{A \multimap B}$, it is easy to check $\pi_\Upsilon(s \cdot m_{\frac{1}{2}}) = \pi_\Upsilon(s) \in \sigma$. So $s \cdot m_{\frac{1}{2}} \in \text{dur}(\sigma)$ □

PROPOSITION J.49.

$$\text{dur}(\text{ccopy}_{A^b}) = \text{durcopy}_A$$

PROOF.

$$\begin{aligned} \text{dur}(\text{ccopy}_{A^b}) &= (K_{\text{Conc}} \text{ccopy}_{A^b})^\# \cap P_{A \multimap A}^{\text{dur}} \\ &= \text{ccopy}_{A^b}^\# \cap P_{A \multimap A}^{\text{dur}} \end{aligned}$$

Now we only need to show $\text{ccopy}_{A^b}^\# \cap P_{A \rightarrow A}^{\text{dur}} = \text{crashcopy}_A \cap P_{A \rightarrow A}^{\text{dur}}$. One direction just follows from J.32. Now we only need to show $\text{crashcopy}_A \cap P_{A \rightarrow A}^{\text{dur}} \subseteq \text{ccopy}_{A^b}^\# \cap P_{A \rightarrow A}^{\text{dur}}$.

Fix $s \in \text{crashcopy}_A \cap P_{A \rightarrow A}^{\text{dur}}$, by definition we know $\pi_\Upsilon(s)$ is well-formed. Furthermore since for each i , $\text{epo}_i(s) \in \text{ccopy}_{A^b}$ so we know $\pi_\Upsilon(s)$ is in ccopy_{A^b} . So $s \in \text{ccopy}_{A^b}^\#$. By assumption, s is durable. So $s \in \text{ccopy}_{A^b}^\# \cap P_{A \rightarrow A}^{\text{dur}}$. \square

PROPOSITION J.50.

$$\text{dur}(-) : \mathbf{Conc} \rightarrow \mathbf{Dur}$$

behaves like an (enriched) lax semifunctor.

PROOF. When $\sigma, \tau \in \mathbf{Conc}$ we have

$$\begin{aligned} \text{dur}(\sigma); \text{dur}(\tau) &= (\sigma^\# \cap P_{A \rightarrow B}^{\text{dur}}); (\tau^\# \cap P_{B \rightarrow C}^{\text{dur}}) \\ &= (\sigma^\#; \tau^\#) \cap P_{A \rightarrow C}^{\text{dur}} \\ &\subseteq (\sigma; \tau)^\# \cap P_{A \rightarrow C}^{\text{dur}} \\ &\subseteq (K_{\mathbf{Conc}}(\sigma; \tau))^\# \cap P_{A \rightarrow C}^{\text{dur}} \\ &= \text{dur}(\sigma; \tau) \end{aligned}$$

\square

J.9 Symmetric Monoidal Structure of Dur

PROPOSITION J.51.

$$\text{Dur} : \mathbf{Conc} \rightarrow \mathbf{Dur}$$

defined by

$$\text{Dur } \sigma := \text{vol}(\sigma) \cap P_{A \rightarrow B}^{\text{dur}}$$

is a functor

PROOF. It follows immediately from the fact that Dur is by definition the composition of two functors. \square

LEMMA J.52. For sets of plays

$$S \subseteq P_{A \rightarrow B} \quad T \subseteq P_{A' \rightarrow B'}$$

and

$$(S \cap P_{A \rightarrow B}^{\text{dur}}) \otimes (T \cap P_{A' \rightarrow B'}^{\text{dur}}) \cap P_{A \otimes A' \rightarrow B \otimes B'}^{\text{dur}} = (S \otimes T) \cap P_{A \otimes A' \rightarrow B \otimes B'}^{\text{dur}}$$

PROOF. Since $S \cap P_{A \rightarrow B}^{\text{dur}} \subseteq S, T \cap P_{A' \rightarrow B'}^{\text{dur}} \subseteq T$, one direction is trivial.

For the other direction, fix $s \in (S \otimes T) \cap P_{A \otimes A' \rightarrow B \otimes B'}^{\text{dur}}$. By definition we know s can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

And for each i , we know

$$s_i \in t_i \otimes u_i$$

and

$$\begin{aligned} t &:= t_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot t_{n+1} \in \sigma \\ u &:= u_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot u_{n+1} \in \tau \end{aligned}$$

Since for each i , $\Upsilon(t_i) \subseteq \Upsilon(s_i)$, and $\Upsilon(u_i) \subseteq \Upsilon(s_i)$, s is durable implies that both t and u are durable.

So $u \in \sigma \cap P_{A \rightarrow B}^{\text{dur}}$, $t \in \tau \cap P_{A' \rightarrow B'}^{\text{dur}}$. So $s \in (S \cap P_{A \rightarrow B}^{\text{dur}}) \otimes (T \cap P_{A' \rightarrow B'}^{\text{dur}}) \cap P_{A \otimes A' \rightarrow B \otimes B'}^{\text{dur}}$ \square

It immediately implies that:

PROPOSITION J.53. *For any given*

$$\sigma : A \multimap B \in \mathbf{Conc} \qquad \sigma' : A' \multimap B' \in \mathbf{Conc}$$

we have

$$\text{Dur}(\sigma \otimes \tau) = \text{Dur} \sigma \boxtimes \text{Dur} \sigma'$$

PROPOSITION J.54. $(\mathbf{Dur}, - \boxtimes -, 1)$ defines a symmetric monoidal category.

PROOF. Functoriality follows from the fact that $- \boxtimes -$ is the composition of two functors.

We start by defining the structural morphisms. The left and right unital are straight-forward. Indeed, they are given by

$$\lambda''_A := \text{Dur} \lambda_A \qquad \rho''_A := \text{Dur} \rho_A$$

where λ and ρ are the left and right unitals in \mathbf{Crash} . The braiding and associator are given b

Now we have enough tools to define the braiding and associator in \mathbf{Dur}

$$\beta''_{A,B} := \text{Dur} \beta_{A,B} \quad \text{and} \quad \alpha''_{A,B,C} := \text{Dur} \alpha_{A,B,C}$$

It immediately follows from functoriality of $- \cap P_{_}^{\text{dur}}$, the definitions of the unitals, associators and braiding, and the fact the corresponding structural morphisms are natural transformations in \mathbf{Crash} that the naturality squares still commute (not that for durable σ , $\sigma \cap P^{\text{dur}} = \sigma$).

The coherence diagrams follow from functoriality of $\text{Dur} -$ together with the fact that $\text{Dur} -$ distributes over $- \boxtimes -$ (Prop. J.53), by noting that all the structural morphisms in \mathbf{Dur} were defined by lifting the corresponding structural morphisms in \mathbf{Crash} , which is also why they are isomorphisms. \square

J.10 Durable Linearizability

PROPOSITION J.55. *For a strategy $v : A^b \in \mathbf{Crash}$ where A is a durable game*

$$\text{dur}(v) = \{s \in P_A^{\text{dur}} \mid s \text{ is durably linearizable with respect to } v\}$$

PROOF. Suppose $s \stackrel{\text{dur}}{\sim} t$ with $t \in v$. So first, by definition, s is durable. Then, $\text{ops}(s) \rightsquigarrow t$, so that $\text{ops}(s) \in K_{\text{Conc}} v$. But then, note that $\pi_{\Gamma}(s) = \text{ops}(s)$ as s is durable so that $s \in (K_{\text{Conc}} v)^{\#} \cap P_A^{\text{dur}} \subseteq \text{dur}(v)$.

For the other direction, suppose $s \in \text{dur}(v)$. Then, $s \in P_A^{\text{dur}}$ and $s \in (K_{\text{Conc}} v)^{\#}$. But then, $\text{ops}(s) = \pi_{\Gamma}(s) \in K_{\text{Conc}} v$, so that $\text{ops}(s) \rightsquigarrow t$ [31], and therefore $s \stackrel{\text{dur}}{\sim} t$. \square

COROLLARY J.56. *For a game $A \in \mathbf{Crash}$, $v' : A \in \mathbf{Crash}$ is durable linearizable to $v : A^b \in \mathbf{Conc}$ if and only if $v' \subseteq \text{dur}(v)$.*

PROPOSITION J.57. *For any durable $\sigma : A \in \mathbf{Conc}$ and $\tau : B \in \mathbf{Conc}$*

$$\text{dur}(\sigma \otimes \tau) = \text{dur}(\sigma) \boxtimes \text{dur}(\tau)$$

PROOF.

$$\begin{aligned}
\text{dur}(\sigma \otimes \tau) &= (K_{\text{Conc}} (\sigma \otimes \tau))^{\#} \cap P_{A \otimes A' \rightarrow B \otimes B'}^{\text{dur}} \\
&= (K_{\text{Conc}} \sigma \otimes K_{\text{Conc}} \tau)^{\#} \cap P_{A \otimes A' \rightarrow B \otimes B'}^{\text{dur}} \\
&= (K_{\text{Conc}} \sigma^{\#} \otimes K_{\text{Conc}} \tau^{\#}) \cap P_{A \otimes A' \rightarrow B \otimes B'}^{\text{dur}} \\
&= (K_{\text{Conc}} \sigma^{\#} \cap P_{A \rightarrow B}^{\text{dur}}) \otimes (K_{\text{Conc}} \tau^{\#} \cap P_{A' \rightarrow B'}^{\text{dur}}) \cap P_{A \otimes A' \rightarrow B \otimes B'}^{\text{dur}} \\
&= \text{dur}(\sigma) \boxtimes \text{dur}(\tau)
\end{aligned}$$

□

PROPOSITION J.58. For $\sigma, \tau : A \multimap B \in \mathbf{Dur}$ and $\sigma', \tau' : A' \multimap B' \in \mathbf{Dur}$ we have

$$\sigma \boxtimes \sigma' \subseteq \tau \boxtimes \tau' \implies \sigma \subseteq \tau \wedge \sigma' \subseteq \tau'$$

PROOF. For given $\sigma, \tau : A \multimap B$ and $\sigma', \tau' : A' \multimap B'$ suppose $\sigma \boxtimes \sigma' \subseteq \tau \boxtimes \tau'$. Now let's try to show $\sigma \subseteq \tau$. Fix $s \in \sigma$, by definition s can be decomposed as

$$s = s_1 \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot s_{n+1}$$

By $\frac{1}{2}$ -receptive we know

$$t := \epsilon \cdot \frac{1}{2} \cdot \dots \cdot \frac{1}{2} \cdot \epsilon \in \sigma'$$

we know that $s \boxtimes t \in \tau \boxtimes \tau'$ but by definition we know it force that $s \in \tau$. So $\sigma \subseteq \tau$

The proof of $\sigma' \subseteq \tau'$ is similar.

□

LOCALITY. For $v'_A : A, v'_B : B \in \mathbf{Dur}$ and $v_A : A, v_B : B \in \mathbf{Conc}$:

$$v'_A \stackrel{\text{dur}}{\sim} v_A \text{ and } v'_B \stackrel{\text{dur}}{\sim} v_B \text{ if and only if } v'_A \boxtimes v'_B \stackrel{\text{dur}}{\sim} v_A \otimes v_B$$

□

PROOF.

$$v'_A \boxtimes v'_B \subseteq \text{dur}(v_A \otimes v_B) = \text{dur}(v_A) \boxtimes \text{dur}(v_B) \iff v'_A \subseteq \text{dur}(v_A) \wedge v'_B \subseteq \text{dur}(v_B)$$

□

PROPOSITION J.59. Let $A, B \in \mathbf{Crash}$. Then $v'_A : A$ is durably linearizable to $v_A : A^b$ if and only if whenever $\sigma : (A \multimap B)^p \in \mathbf{Conc}$ implements a concurrent object linearizable to v_B using v_A , then $\text{dur}(\sigma) : A \multimap B$ implements an object durably linearizable to v_B using v'_A .

PROOF. For the forward direction, we have that by assumption

$$v_A; \sigma \subseteq K_{\text{Conc}} v_B$$

And by lax functoriality of $\text{dur}(-)$

$$v'_A; \text{dur}(\sigma) \subseteq \text{dur}(v_A); \text{dur}(\sigma) \subseteq \text{dur}(v_A; \sigma) \subseteq \text{dur}(K_{\text{Conc}} v_B) = \text{dur}(v_B)$$

For the backward direction, note that

$$v_A; \text{ccopy}_{A^b} = K_{\text{Conc}} v_A$$

So, by assumption,

$$v'_A; \text{dur}(\text{ccopy}_{A^b}) \subseteq \text{dur}(K_{\text{Conc}} v_A)$$

and hence,

$$v'_A = v'_A; \text{durcopy}_A = v'_A; \text{dur}(\text{ccopy}_{A^b}) \subseteq \text{dur}(K_{\text{Conc}} v_A) = \text{dur}(v_A)$$

□