



Modal Crash Types for Intermittent Computing^{*}

Farzaneh Derakhshan^(✉), Myra Dotzel, Milijana Surbatovich, and Limin Jia

Carnegie Mellon University, Pittsburgh PA, USA
{fderakhs,mdotzel,milijans,liminjia}@andrew.cmu.edu

Abstract. Intermittent computing is gaining traction in application domains such as Energy Harvesting Devices (EHDs) that experience arbitrary power failures during program execution. To make progress, programs require system support to checkpoint state and re-execute after power failure by restoring the last saved state. This re-execution should be *correct*, i.e., simulated by a continuously-powered execution. We study the logical underpinning of intermittent computing and model checkpoint, crash, restore, and re-execution operations as computation on Crash types. We draw inspiration from adjoint logic and define Crash types by introducing two adjoint modality operators to model persistent and transient memory values of partial (re-)executions and the transitions between them caused by checkpoints and restoration. We define a Crash type system for a core calculus. We prove the correctness of intermittent systems by defining a novel logical relation for Crash types.

Keywords: intermittent computing · modal Crash type · logical relation

1 Introduction

Intermittent computing is gaining importance in application domains that require inaccessible or large-scale device deployments, such as wildlife monitoring [28], tiny satellites [22,29], or smart civil infrastructure [1]. As battery maintenance may be infeasible in these environments, programs can instead run on batteryless Energy Harvesting Devices (EHDs). An EHD can run solely off energy harvested from its environment, at the cost of being powered intermittently. The device harvests energy (e.g., via solar panel) into a re-chargeable buffer. Once the energy buffer is full, the device turns on and begin to compute, consuming the stored energy. When the buffer drains, the device turns off at an arbitrary location until it can recharge and repeat this operational cycle. A power failure erases volatile execution state (e.g., the program counter), while

* This work was generously funded in part through National Science Foundation (NSF) Award 2007998, NSF Graduate Research Fellowship Program grants DGE1745016 and DGE2140739, and the CMU CyLab Security & Privacy Institute. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsoring organizations.

nonvolatile state persists. For programs to make progress, they require *intermittent system* support to save state at checkpoints and restore the saved state after power failure, potentially causing re-execution from the last checkpoint.

As EHDs aim to enable long-term deployments with little or no maintenance, intermittent systems must execute programs reliably despite frequent power failures and partial executions. Initial systems [35,43,24] relied only on informal notions of correctness that left them susceptible to memory consistency bugs caused by reading the results of partial executions [23] or by allowing sensor reads from past executions to remain in the nonvolatile memory [39]. More recent work [41,40,9,13] provides formal frameworks and correctness criteria for reasoning about intermittent execution. More concretely, all intermittent executions of a program must be simulated by some continuously-powered execution [41]. In other words, intermittent execution should be *idempotent*. Even if the system induces multiple partial executions of a program due to power failure, the program should not generate a different result than it would on a single execution.

The correctness of an intermittent execution relies on checkpointing, restoring, and finalizing state upon reaching the next checkpoint; mistakes in these operations can lead to incorrect, non-idempotent behavior. Few works have tried to understand the *fundamental logical underpinning* of these operations. This work fills this gap by formalizing checkpointing, crash, restoration, and re-execution as computation on *Crash types*. Crash types capture the core notion of intermittent computing: some values and computations persist across power failures and others do not. For instance, nonvolatile memory state persists across power failure and reboots, while volatile memory does not. Conversely, partially computed results do (or rather *should*) not persist across power failures, while completed/checkpointed computations do. We call the former *unstable* values and computations and the latter *stable* values and computations. Our key insight is that the interactions between these stable and unstable components bear close resemblance to shifts in adjoint logic [8,36]. Computation of a stable value can only rely on locations that store stable values, while computation on unstable values can rely on both stable and unstable values. Moreover, checkpoint and restore operations can turn values of one type to the other. We define terms and their associated types so that each of the key intermittent computing operations must be well-typed under our Crash types.

We define a core calculus for intermittent computing and develop a type system for Crash types by using the two adjoint modality operators. The Crash type of an intermittent computation is: $\mathbf{C}_{\text{unit}} = \downarrow(\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}) \vee \downarrow \uparrow \text{unit}$, which says that the computation will either encounter a power failure (the left disjunct), or succeed in producing a stable value (the right disjunct). In the former case, the computation is suspended until energy arrives, after which it will again act as an intermittent computation. This recursive definition captures the multiple re-executions of a computation under repeated power failures. To prove the correctness of intermittent systems, we define a novel logical relation for Crash types, indexed by the number of power failures, which relates a continuously-

powered execution to an intermittent execution. While intermittent computing motivates our results, the methods we develop are generally applicable to other system failures with the same effect on persistent and transient storage.

This paper makes the following technical contributions:

- The first logical interpretation of key operations of intermittent execution.
- Novel Crash types to specify how stable and unstable portions of the system and computation interact.
- A core calculus for Crash types with progress and preservation.
- A novel logical relation to prove the correctness of intermittent executions.

Detailed proofs and definitions can be found in the extended TR [15].

2 Background

We provide background on intermittent computing and detail how checkpoint systems work to store and restore program state to handle power failures.

Intermittent Computing on EHDs. EHDs need intermittent system support to save necessary state before power failure and to restore it after reboot. When and where such checkpoints occur governs the *intermittent execution model* under which software executes. The two prevailing intermittent execution models are just-in-time (JIT) checkpoints [5,4] and atomic execution [23,24,43,37]. Under a JIT model, state is saved immediately prior to power failure so that execution resumes from the same point after reboot. Under an atomic execution model, state is saved at the beginning of an *atomic region*. If power fails before the end of the region, the system will reboot to the beginning of the region, re-executing until the region completes without power failure (akin to software transactions [38]). State-of-the-art intermittent systems use a hybrid “JIT + Atomics” model that defaults to JIT checkpoints except when there is an explicit atomic region [40,25,19]. Our core calculus follows this hybrid model.

To ensure idempotence, an intermittent system must save the value of volatile state and often a portion of the *nonvolatile* state. To illustrate why, consider an execution of the simple program in Fig. 1. The program has four variables stored in nonvolatile memory: x , y , and z of type `int` and u of type `bool`. It consists of two code blocks: an atomic region declared with the `Ckpt` construct (lines 1-7 on the left of Fig. 1) and a regular code block executed in JIT mode (lines 8-14 on the right). A continuous execution of the atomic region with initial state $x = 2, y = 0, z = 1, u = \text{ff}$ ends in $x = 2, y = 1, z = 1, u = \text{tt}$. Now, suppose power fails after the execution of Line 2. Once the device recharges, the program restarts from the start of the atomic region. If the system does not restore y 's original value, this re-run computes an incorrect result: $x = 2, y = 2, z = 1, u = \text{ff}$. Thus, to ensure idempotent execution, an intermittent system must checkpoint, i.e., save the value of, both volatile and nonvolatile memory. We next explain correct execution of the program in Fig. 1 for atomic and JIT modes.

Atomic Region Execution. As EHDs are highly resource constrained, the system should save state judiciously; checkpointing all of nonvolatile memory is

```

1      Ckpt[a1; x,z:read-only](      8      let w=not u in
2          y:=y+z;                    9          if w then
3          let w= x-y in              10             x=x+y;
4              if w>0 then            11             w=ff
5                  u:=tt              12             else
6              else                    13             skip;
7          u:=ff);                    14      skip
    
```

Fig. 1. An example program with an atomic region and a JIT region

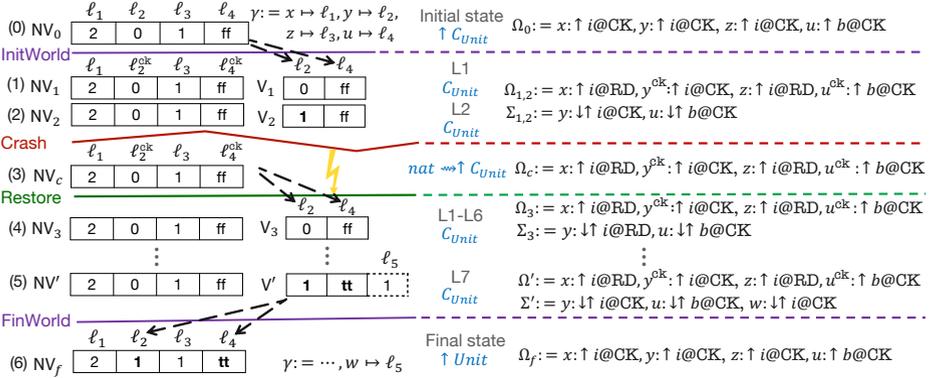


Fig. 2. Intermittent execution of an atomic region. We write i for `int` and b for `bool`.

expensive and unnecessary. For example, variables in an atomic region that are read-only (i.e., never updated) do not change value and need not be checkpointed. In our example, x and z are read-only, so checkpointing y and u is enough to ensure correct intermittent execution. Many intermittent systems follow this design of checkpointing all variables that are not read-only [37,19,17,26,44,12]. Given such a system, Fig. 2 shows an execution of the atomic region in Fig. 1. For now, ignore the last two columns about typing. To save and restore state, the system follows redo-log semantics. It records updates to checkpointed variables in a special volatile region, not main memory. This region clears if power fails, throwing out partial updates. Upon reaching the next atomic or JIT region, the system commits the updates by copying them back to main memory.

Row (0) shows initial nonvolatile locations, their values, and the mapping between variables and memory locations; locations ℓ_1, ℓ_2, ℓ_3 , and ℓ_4 in the non-volatile memory correspond to variables x, y, z and u , respectively. When starting the atomic region (Row (1)), the system takes a snapshot of ℓ_2 and ℓ_4 and stores it in the volatile region V_1 . We mark the original nonvolatile locations as checkpointed with the superscript `ck`, i.e., ℓ_2^{ck} and ℓ_4^{ck} . Checkpointed locations ℓ_2^{ck} and ℓ_4^{ck} remain untouched for the remainder of the atomic region execution. Every access to variables y and u will instead be associated with their volatile copy ℓ_2 and ℓ_4 , e.g., the assignment in Line 2 is applied to the volatile logs of Row (2).

On power failure, all volatile memory clears (Row (3)), throwing out the log. The system shuts down until more energy is harvested, at which point the system regenerates the volatile copies ℓ_2 and ℓ_4 (Row (4)) and resumes execution from Line 2. When the execution of the atomic region is complete (Row (5)), the system commits the updated values of the checkpointed locations (ℓ_2 and ℓ_4) from volatile memory to their original nonvolatile locations (Row (6)). During execution, local variables are stored to volatile memory via a `let` construct, e.g., location ℓ_5 for variable w on Line 3, corresponding to a volatile execution stack. On power failure, the device clears all volatile memory, but such stack allocated locations will be recreated upon re-execution.

JIT Region Execution. The JIT execution model prevents re-execution, so the intermittent system only saves and restores volatile state at checkpoints. Fig. 3 shows the details of executing the code on the right of Fig. 1 in JIT mode. Row (0) shows the initial nonvolatile locations, their values, and the mapping from variables to locations. The system starts the JIT region by creating an empty context to be populated by volatile locations (Row (1)). The `let` construct in Line 8 allocates a fresh location ℓ_5 in V_2 and updates the mapping to associate variable w to ℓ_5 . On a power failure in JIT mode, the system creates a nonvolatile copy of the volatile location ℓ_5 just before it loses the location (Row (3)). It marks the nonvolatile copy with the superscript `ck`. When restoring the program, the system restores these copies to volatile memory and dismisses the nonvolatile backups (Row (4)). The program then continues with the `if` clause on lines 9-12, finally dropping the volatile location ℓ_5 , as it is out of scope (Row (5)).

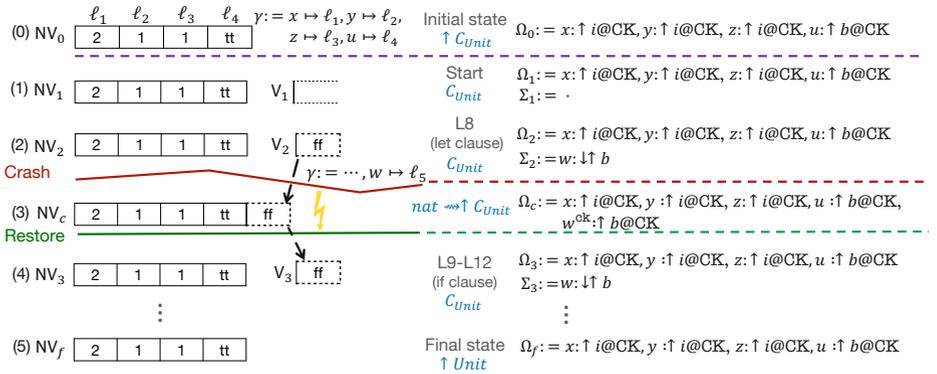


Fig. 3. Intermittent execution of a JIT region. We write i for `int` and b for `bool`.

3 Key Ideas of Crash Types

We present the intuition behind the stable and unstable memory types (Sec. 3.1), Crash types which internalize checkpointing, power failure/crash, restoration, re-

execution, and finalization of atomic regions (Sec. 3.2), and the independence principle applied to intermittent computing (Sec. 3.3).

3.1 Modal Store Types

An unstable value is an intermediate result of an execution towards a stable value and will be lost upon a power failure. However, if the result of a partial execution is committed to a nonvolatile location, it will persist and is thus stable. To reflect the behavior of a memory location in its type, we introduce two (adjoint) modalities \uparrow_u^s (read as “up shift from unstable to stable”) and \downarrow_u^s (read as “down shift from stable to unstable”), where $\uparrow_u^s \tau$ indicates that the location stores a stable value of type τ and $\downarrow_u^s \tau$ indicates that the location stores an intermediate result of an execution toward a value of type τ . To fully capture how intermittent execution interacts with a memory location, we also annotate the type of a memory location with an access qualifier, RD or CK, that represents whether the location is read-only or checkpointed by the system, respectively.

In our example in Fig. 2, the read-only variable x is stored in nonvolatile memory, so it has type $x : \uparrow_u^s \text{int@RD}$. The checkpointed variable y has type $y^{\text{ck}} : \uparrow_u^s \text{int@CK}$ in the nonvolatile memory, while y ’s volatile copy has type $y : \downarrow_u^s \uparrow_u^s \text{int@CK}$. We use the context Ω to type nonvolatile memory and the context Σ to type volatile memory, as shown in the third columns of Figs. 2 and 3. We drop the superscript s and subscript u from the modalities for brevity.

3.2 Crash Types

To capture the effects of intermittent execution in the type of expressions and commands, we introduce *Crash types*, as the notion of stable and unstable values is insufficient. One might expect the expression $x - y$ to have the type $\downarrow \uparrow \text{int}$ as it is a (partial) execution towards computing a stable integer value. However, this type does not account for steps due to power failure: the crash itself, waiting for the device to charge, restoration, and re-execution. To reflect these runtime system steps at the type level, we assign the expression a type in the form of a disjunction $\boxed{?} \vee \downarrow \uparrow \text{int}$, where $\boxed{?}$ is a type for computations that handle power failures. This type means that the expression either power fails, or completes its execution that evaluates to int . Next, we fill in $\boxed{?}$ for commands and expressions. $\boxed{?}$ is a recursive type since it handles re-execution.

Commands. The Crash type for commands is: $\mathbf{C}_{\text{unit}} = \downarrow(\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}) \vee \downarrow \uparrow \text{unit}$. The right disjunct states that if no power failure occurs while executing a command, then it computes a stable value of type unit . The left disjunct states that on power failure, the computation continues as a function; after receiving a (logical) energy input from the environment, it becomes a computation that yields a stable value of a command type, i.e., \mathbf{C}_{unit} . This computation will execute after the restore, which differs for atomic and JIT modes. In an atomic region, the system re-executes the region from the beginning, and in a JIT region, the system continues with the same command that was interrupted by the failure.

Expressions. The definition of the Crash type for expressions depends on the execution mode, just as the continuation of the program after a power failure depends on the mode. In an atomic region, the system restores an interrupted run of the expression to the original command enclosed in the region, so the type of an atomic mode expression is $\mathbf{C}_A^{\text{atom}} = \downarrow(\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}) \vee \downarrow \uparrow A$, where the left disjunct is the same as that of a command. On the other hand, an interrupted run of an expression in JIT mode will be restored to the expression itself. Hence, the type of a JIT mode expression is $\mathbf{C}_A^{\text{jit}} = \downarrow(\text{nat} \rightsquigarrow \uparrow \mathbf{C}_A^{\text{jit}}) \vee \downarrow \uparrow A$, where the left disjunct states that after power failure and reception of the energy input, the computation again yields a stable value of a JIT mode expression type.

3.3 Independence Principle for Typing Intermittent Execution

We design our typing rules to follow the rules for \downarrow and \uparrow modalities in adjoint logic. We introduce two judgment categories. The first category (J_s) is for deriving stable types and corresponds to the judgments of the form $\Omega \vdash \tau^s$, meaning that the rules can rely only on stable locations to evaluate computation on a stable type. The second category (J_u) is for deriving unstable types and corresponds to the judgments of form $\Omega; \Sigma \vdash \tau^u$, meaning that the rules can rely on both stable and unstable locations to evaluate computation on an unstable type.

The adjoint modalities allow going back and forth between judgments J_s and J_u , mirroring checkpointing and restoration operations. The following four sequent calculus rules in the underlying logic govern this back-and-forth behavior in our system. The rules are derivable from the more general rules in prior work [8,34,36]—in particular, the $\uparrow L^*$ rule can be derived from a cut rule and $\downarrow L$. Typical of sequent calculus style rules, we read them bottom-up and match each execution step of a command with the reading of a corresponding rule. Next, we illustrate this matching using the execution steps in Figs. 2 and 3.

$$\frac{\Omega; \cdot \vdash \tau^u}{\Omega \vdash \uparrow \tau^u} \uparrow R \quad \frac{\Omega, \uparrow A^u; \Sigma, \downarrow \uparrow A^u \vdash \tau^u}{\Omega, \uparrow A^u; \Sigma \vdash \tau^u} \uparrow L^* \quad \frac{\Omega \vdash \tau^s}{\Omega; \Sigma \vdash \downarrow \tau^s} \downarrow R \quad \frac{\Omega, \uparrow A^u; \Sigma \vdash \tau^u}{\Omega; \Sigma, \downarrow \uparrow A^u \vdash \tau^u} \downarrow L$$

Shifts in Atomic Mode (Fig. 2): A combination of $\uparrow R$ and two $\uparrow L^*$ rules corresponds to creating a volatile log from the nonvolatile locations when starting the atomic region, i.e., the step from Row (0) to Row (1). The last two columns in Row (0) correspond to the conclusion of a $\uparrow R$ rule: $\Omega_0 \vdash \uparrow \mathbf{C}_{\text{unit}}$. An application of $\uparrow R$ from bottom to top drops the \uparrow modality from the type of the program and opens an empty volatile region, i.e., $\Omega_0; \cdot \vdash \mathbf{C}_{\text{unit}}$. Next, one application of $\uparrow L^*$, copies the variable y of type $\uparrow \text{int}$ to the volatile memory with the type $\downarrow \uparrow \text{int}$. Similarly, the next application of $\uparrow L^*$ copies the variable u of type $\uparrow \text{bool}$ to the volatile memory with the type $\downarrow \uparrow \text{bool}$. The same combination corresponds to creating a volatile log from a nonvolatile location when restarting the atomic region, i.e., the step from Row (3) to Row (4), again copying variables y and u to the volatile memory.

The $\downarrow R$ rule corresponds to a power failure, which erases the volatile memory Σ . From Row (2) to Row (3) in Fig. 2, the system loses the volatile locations of y

and u and closes off the volatile context. Row (2) corresponds to the conclusion of the rule, and Row (3) corresponds to its premise. The type of the command in Row (2) changes from C_{unit} to $\downarrow(\text{nat} \rightsquigarrow \uparrow C_{\text{unit}})$ (by another \vee -R rule as a crash is detected), and then to the type $(\text{nat} \rightsquigarrow \uparrow C_{\text{unit}})$ in Row (3).

Finally, a $\downarrow L$ rule combined with a standard weakening rule and a $\downarrow R$ rule corresponds to the final commit of the volatile context, i.e., stepping from Row (5) to Row (6), the nonvolatile context drops the locations y and u of types $\uparrow \text{int}$ and $\uparrow \text{bool}$, respectively, by a weakening rule. These two variables map to the locations with outdated values. Next, the volatile locations of y and u in Σ' , which contain the up-to-date values, commit their values to the nonvolatile context by a $\downarrow L$ rule. Then, a $\downarrow R$ rule closes off the remaining volatile context, which contains w of type $\downarrow \uparrow \text{int}$. The type of the command in Row (2) changes from C_{unit} to $\downarrow \uparrow \text{unit}$ (by a separate \vee -R rule as the system detects a successful execution) and from that to type $\uparrow \text{int}$ in Row (6).

Shifts in JIT Mode (Fig. 3): A $\uparrow R$ rule corresponds to creating an empty volatile context Σ_1 when starting the JIT region, i.e., the step from Row (0) to Row (1). A combination of the $\downarrow L$ rule and $\downarrow R$ rule corresponds to a power failure, i.e., the stepping from Row (2) to Row (3). A $\downarrow L$ rule copies the location w of type $\downarrow \uparrow \text{bool}$ from volatile memory Σ_2 to nonvolatile memory Ω_c . A $\downarrow R$ rule closes off the (empty) nonvolatile memory. As in atomic mode, a combination of $\uparrow R$ and $\uparrow L^*$ rules corresponds to creating a volatile log from a nonvolatile location when restarting the command after the failure, i.e., the step from Row (3) to Row (4). The $\uparrow R$ rule clears a portion of volatile memory, and the $\uparrow L^*$ rule copies variable w from nonvolatile memory into volatile memory. We need an extra weakening rule to eliminate the remaining variable w in nonvolatile memory. The dropping of volatile memory at the end of execution (Row (5)) is not a modal step, but rather follows from a standard rule for the let clause.

4 A Basic Calculus for Intermittent Execution

We present the syntax, semantics, and the Crash type system for a basic calculus.

4.1 Syntax

The syntactic constructs are summarized in Fig. 4. Expressions include constants, variables, and binary operations while commands include assignments, mutable let bindings, sequencing, and if branching. A program consists of sequenced blocks of commands and atomic regions, denoted $\text{Ckpt}[\text{aID}, \rho](c)$ with a unique identifier aID , read-only variables ρ , and the enclosed command c .

Nonvolatile memory (NV) and volatile memory (V) map locations ℓ to values. Each location is annotated with its access mode q (RD or CK). The nonvolatile memory location ℓ_{ck} is the checkpointed copy of location ℓ in volatile memory. The context γ maps variable names to memory locations. Access mode qualifiers in V and NV have constrained values (to be discussed in the semantics).

Command, expression, and memory

<i>values</i> $v ::= n \mid \text{tt} \mid \text{ff} \mid x$ <i>exprs</i> $e ::= v \mid e \odot e$ <i>cmds</i> $c ::= \text{skip} \mid \text{let } x = e \text{ in } c \mid c; c$ $\mid \text{if } e \text{ then } c \text{ else } c \mid x ::= e$ <i>progs</i> $p ::= \text{Ckpt}[\text{aID}, \rho](c); p \mid c; p \mid \text{skip}$	<i>access qualifier</i> $q ::= \text{CK} \mid \text{RD}$ <i>var loc map</i> $\gamma ::= \cdot \mid \gamma, x \mapsto \ell$ <i>nonvolatile mem</i> $\text{NV} ::= \cdot \mid \ell @ q \leftrightarrow v, \text{NV}$ $\mid \ell_{\text{ck}} @ \text{CK} \leftrightarrow v, \text{NV}$ <i>volatile mem</i> $\text{V} ::= \cdot \mid \ell @ \text{CK} \leftrightarrow v, \text{V}$
--	--

Instructions, statements, and configurations.

<i>commands</i> $c ::= \dots c;_W c$ <i>continuations</i> $\kappa ::= c \mid e$ <i>statements</i> $s ::= \kappa \mid i \mid p$ <i>energy level</i> $g ::= \cdot \mid n$ <i>charge stream</i> $\chi ::= n :: \chi$ <i>exec. mode</i> $\text{Md} ::= \text{aID}(c) \mid \text{jit}$	<i>crash instrs</i> $i ::= \downarrow \varepsilon \# \text{in}(b > 0, \uparrow \kappa)$ $\mid \varepsilon \# \text{in}(b > 0, \uparrow \kappa) \mid \uparrow \kappa$ <i>open config</i> $K_o ::= (\gamma \mid \text{Md} \mid g \mid \text{NV} \mid \text{V} \mid s)$ $\mid (\gamma \mid \text{Md} \mid g \mid \text{NV} \mid s)$ <i>closed config</i> $K_c ::= [\chi \triangleright \varepsilon] \otimes K_o$
--	---

Fig. 4. Summary of syntax

The runtime instruction $c_1;_W c_2$ is used for evaluating c_1 under the execution context W . To model energy harvesting from the environment, we assume a unique external energy channel, ε , from which the system receives energy. Three crash instructions control the system in the event of a power failure. The instruction $\downarrow \varepsilon \# \text{in}(b > 0, \uparrow \kappa)$ models the system that faces a power failure, where κ is the interrupted command or expression, and $b > 0$ is a guard to ensure that the bound incoming energy variable b is positive. The instruction $\varepsilon \# \text{in}(b > 0, \uparrow \kappa)$ models the system awaiting an energy input to be bound to b . The instruction $\uparrow \kappa$ models the system ready to restore memory and re-execute.

We write K_o to denote an *open* system configuration, consisting of the mapping γ , the mode of execution Md (i.e., atomic or JIT), energy available for this execution g , memories, and the statement s to be executed. The energy level (\cdot) models the state right after power failure. We close an open configuration with $[\chi \triangleright \varepsilon]$; we connect it via an external energy channel ε to an infinite charging stream Ξ of natural numbers, which models available energy the configuration harvests from the environment at each power failure point for re-execution.

We call a configuration that cannot take a step a value configuration (value for short). An open configuration of form $(\dots \mid g \mid \dots \mid s)$ is a value, i.e., $\text{Val}(\dots \mid g \mid \dots \mid s)$, if either s is a constant or skip , it has depleted all energy for this execution ($g = 0$), or s is a crash instruction. The latter two cases are values because they cannot take a step without interacting with the environment or perform operations on the volatile and nonvolatile memory specific to handling power failures. A closed configuration is a value only if the statement s is skip with some energy left ($g > 0$). We list all values in the extended TR [15].

4.2 Operational Semantics

Top-level Program Execution. The top-level semantic rules for setting up and finalizing the atomic and JIT execution contexts are shown in Fig. 5. The P-CKPT rule applies if the next code block is an atomic region. The nonvolatile

$$\begin{array}{c}
 n > 0 \quad \text{InitWorld}_d(\text{NV}; \rho; \gamma) = \text{NV}_0, V_0 \\
 [\chi \triangleright \varepsilon] \otimes \gamma \mid \text{aID}(c_0) \mid n \mid \text{NV}_0 \mid V_0 \mid c_0 \Rightarrow^* [\chi' \triangleright \varepsilon] \otimes \gamma' \mid \text{aID}(c_0) \mid n' \mid \text{NV}' \mid V' \mid \text{skip} \\
 n' > 0 \quad \text{NV}_1 = \text{FinWorld}_d(\text{NV}'; V') \\
 \hline
 [\chi \triangleright \varepsilon] \otimes \gamma \mid n \mid \text{NV} \mid \text{Ckpt}[(\text{aID}; \rho)](c_0); p \Rightarrow [\chi' \triangleright \varepsilon] \otimes \gamma' \mid n' \mid \text{NV}_1 \mid p \quad (\text{P-CKPT})
 \end{array}$$

$$\begin{array}{c}
 n > 0 \quad n' > 0 \\
 [\chi \triangleright \varepsilon] \otimes \gamma \mid \text{jit} \mid n \mid \text{NV} \mid \cdot \mid c \Rightarrow^* [\chi' \triangleright \varepsilon] \otimes \gamma' \mid \text{jit} \mid n' \mid \text{NV}' \mid V' \mid \text{skip} \\
 \hline
 [\chi \triangleright \varepsilon] \otimes \gamma \mid n \mid \text{NV} \mid c; p \Rightarrow [\chi' \triangleright \varepsilon] \otimes \gamma' \mid n' \mid \text{NV}' \mid p \quad (\text{P-SEQ})
 \end{array}$$

Fig. 5. Closed configuration semantics for programs

NV_0 and volatile V_0 locations are initialized based on a given NV , declared read-only variables ρ , and their mapping γ to locations. The InitWorld_d function (a) changes the qualifier of locations in NV that are declared as read-only in ρ from CK to RD , (b) creates V_0 by copying the rest of the locations of NV that still have qualifier CK , and (c) marks the original version of the locations ℓ in NV that still have qualifier CK as checkpointed (ℓ_{ck}). This part corresponds to the step from Row (0) to Row (1) in Fig. 2. The closed configuration of c_0 is evaluated until completion, using the rules in Fig. 6. This execution may undergo several power failures and corresponds to the steps from Row (1) to Row (5) in Fig. 2. Finally, the FinWorld_d function closes off atomic regions, finalizing the volatile and nonvolatile locations. FinWorld_d (a) copies the values of volatile locations in V' that have a checkpointed version into NV' , (b) removes CK from the locations in NV' , i.e., converts ℓ_{ck} to ℓ , and (c) replaces the RD qualifier of the locations in NV' with CK . This corresponds to the step from Row (5) to Row (6) in Fig. 2.

The P-SEQ rule applies when the next code block is a regular command c . The closed configuration of c with an empty initial set of volatile locations is fully evaluated. This corresponds to the steps from Row (0) to Row (1) and Row (1) to Row (5) in Fig. 3. Then the resulting volatile locations V' scoped in c are dropped, corresponding to the step from Row (5) to Row (6) in Fig. 3.

Command Execution (Closed Config). We summarize rules for a closed configuration in the top part of Fig. 6. Rule D-STEP steps the closed command configuration when the corresponding open configuration steps. Next, we explain the trio of power failure, charge, and restore rules. When the energy for this execution is depleted (i.e., $g = 0$), the D-CRASH rule applies, stepping the system to the crash instruction $\downarrow \varepsilon \# \text{in}(b > 0; \uparrow \kappa)$. Next, D-S-JIT or D-S-AID rules apply and operate on volatile memory based on the execution mode Md . In JIT mode, D-S-JIT checkpoints and stores all volatile memory in nonvolatile locations. In atomic mode, D-S-AID drops all volatile memory locations. Then, D-CHARGE applies and inputs a natural number $n > 0$ from the energy channel, replenishing the configuration's energy level for re-execution. Finally, the program is restored via D-RESTORE-JIT and D-RESTORE-AID which copy checkpointed locations into volatile memory. D-RESTORE-JIT drops the checkpointed regions and steps

Closed Configuration Semantics for Commands and Crash Instructions

$$\begin{array}{c}
\frac{\gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \mathbf{V} \mid c \rightarrow \gamma \mid \mathbf{Md} \mid n' \mid \mathbf{NV}' \mid \mathbf{V}' \mid c'}{[\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \mathbf{V} \mid c \Rightarrow [\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid n' \mid \mathbf{NV}' \mid \mathbf{V}' \mid c'} \text{ (D-STEP)} \\
\hline
[\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid 0 \mid \mathbf{NV} \mid \mathbf{V} \mid c \Rightarrow [\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid \cdot \mid \mathbf{NV} \mid \mathbf{V} \mid \downarrow \varepsilon \# \text{in}(b > 0; \uparrow c) \text{ (D-CRASH)} \\
\hline
\frac{\mathbf{Md} = \text{jit}}{[\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid \cdot \mid \mathbf{NV} \mid \mathbf{V} \mid \downarrow \varepsilon \# \text{in}(b > 0; \uparrow \kappa) \Rightarrow [\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid \mathbf{NV}, \mathbf{V}_{\text{ck}} \mid \varepsilon \# \text{in}(b > 0; \uparrow \kappa)} \text{ (D-S-JIT)} \\
\hline
\frac{\mathbf{Md} = \text{aID}(c_0) \quad \gamma' \subseteq \gamma \quad \text{range}(\gamma') = \text{dom}(\mathbf{NV})}{[\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid \cdot \mid \mathbf{NV} \mid \mathbf{V} \mid \downarrow \varepsilon \# \text{in}(b > 0; \uparrow \kappa) \Rightarrow [\chi \triangleright \varepsilon] \otimes \gamma' \mid \mathbf{Md} \mid \cdot \mid \mathbf{NV} \mid \varepsilon \# \text{in}(b > 0; \uparrow \kappa)} \text{ (D-S-AID)} \\
\hline
[n :: \chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid \cdot \mid \mathbf{NV} \mid \varepsilon \# \text{in}(b > 0; \uparrow \kappa) \Rightarrow [\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \uparrow \kappa \text{ (D-CHARGE)} \\
\hline
\frac{\mathbf{NV} = \mathbf{NV}', \mathbf{NV}''_{\text{ck}}}{[\chi \triangleright \varepsilon] \otimes \gamma \mid \text{jit} \mid n \mid \mathbf{NV} \mid \uparrow \kappa \Rightarrow [\chi \triangleright \varepsilon] \otimes \gamma \mid \text{jit} \mid n \mid \mathbf{NV}' \mid \mathbf{NV}'' \mid \kappa} \text{ (D-RESTORE-JIT)} \\
\hline
\frac{\mathbf{NV} = \mathbf{NV}', \mathbf{NV}''_{\text{ck}}}{[\chi \triangleright \varepsilon] \otimes \gamma \mid \text{aID}(c_0) \mid n \mid \mathbf{NV} \mid \uparrow \kappa \Rightarrow [\chi \triangleright \varepsilon] \otimes \gamma \mid \text{aID}(c_0) \mid n \mid \mathbf{NV} \mid \mathbf{NV}'' \mid c_0} \text{ (D-RESTORE-AID)} \\
\hline
\text{Selected expression and command semantics} \\
\hline
\frac{\gamma = \gamma', [x \mapsto \ell] \quad \mathbf{V} = \ell @ q \hookrightarrow v, \mathbf{V}' \quad n = n' + 1}{\gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \mathbf{V} \mid x \rightarrow \gamma \mid \mathbf{Md} \mid n' \mid \mathbf{NV} \mid \mathbf{V} \mid v} \text{ (D-V-READ)} \\
\hline
\frac{\text{Val}(\gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \mathbf{V} \mid e) \quad \mathbf{V} = \mathbf{V}', \ell @ q \hookrightarrow v' \quad q \neq \text{RD} \quad \gamma = \gamma', [x \mapsto \ell] \quad n = n' + 1}{\gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \mathbf{V} \mid x := e \rightarrow \gamma \mid \mathbf{Md} \mid n' \mid \mathbf{NV} \mid \mathbf{V}', \ell @ q \hookrightarrow e \mid \text{skip}} \text{ (D-ASSIGN-V)}
\end{array}$$

Fig. 6. Statement steps

to the interrupted command κ , while D-RESTORE-AID keeps the checkpointed regions and steps to the original command c_0 in the atomic region.

Command/Expression Execution (Open Config). The rules for executing commands and expressions in an open configuration are standard. We present a selection of them on the bottom of Fig. 6. Each step decrements the energy level by one. The rules ensure that checkpointed location ℓ_{ck} in \mathbf{NV} is not read by the program, as it could store outdated data, and is not written to, as this would tamper with the checkpointed value.

4.3 Types, Typing Contexts, and Judgments

This section introduces the typing judgments used in our static typing.

(J_u)	$\text{Md} \mid b \mathcal{R} 0 : \text{nat} \mid \Omega; \Sigma \vdash c :: \mathbb{C}_{\text{unit}}$	c could crash
(J_u)	$\text{Md} \mid b : \text{nat} \mid \Omega; \Sigma \vdash \text{skip} :: \downarrow \uparrow \text{unit}$	c will not crash
(J_s)	$\text{Md} \mid b : \text{nat} \mid \Omega \vdash \text{skip} :: \uparrow \text{unit}$	after commit
(J_u)	$\text{Md} \mid b \mathcal{R} 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{RD}} e :: \mathbb{C}_A^{\text{Md}}$	e read, could crash
(J_s)	$\text{Md} \mid b : \text{nat} \mid \Omega; \Sigma \vdash_{\text{RD}} v :: \downarrow \uparrow A$	e read no crash
(J_s)	$\text{Md} \mid b : \text{nat} \mid \Omega; \Sigma \vdash_{\text{RD}} v :: \uparrow A$	e read, commit
(J_u)	$\text{Md} \mid b : \text{nat} \mid \Omega; \Sigma \vdash_{\text{WT}} x :: \downarrow \uparrow A$	write on x , no crash
(J_s)	$\text{Md} \mid b : \text{nat} \mid \Omega \vdash_{\text{WT}} x :: \uparrow A$	write on x , commit
(J_s)	$\text{Md} \mid b : \text{nat} \mid \Omega \vdash p :: \uparrow \mathbb{C}_{\text{unit}}$	before execution
(J_u)	$\text{Md} \mid b = 0 : \text{nat} \mid \Omega; \Sigma \vdash \kappa :: \mathbb{C}_T^{\text{Md}}$	about to crash
(J_u)	$\text{Md} \mid \cdot \mid \Omega; \Sigma \vdash \downarrow \varepsilon \# \text{in}(b > 0, \uparrow \kappa) :: \downarrow (\text{nat} \rightsquigarrow \uparrow \mathbb{C}_T^{\text{Md}})$	crash state
(J_s)	$\text{Md} \mid \cdot \mid \Omega \vdash \varepsilon \# \text{in}(b > 0, \uparrow \kappa) :: \text{nat} \rightsquigarrow \uparrow \mathbb{C}_T^{\text{Md}}$	waiting for energy
(J_s)	$\text{Md} \mid b > 0 : \text{nat} \mid \Omega \vdash \uparrow \kappa :: \uparrow \mathbb{C}_T^{\text{Md}}$	before re-execution

Table 1. Typing judgment summary

Types and Static Context. Our types are summarized below. The two modalities stratify types into the varieties stable (τ^s) and unstable (τ^u). The base store types `int` and `bool` are considered unstable. A type variable v_t denotes a type in the set $\{\mathbb{C}_{\text{unit}}, \mathbb{C}_A^{\text{atom}}, \mathbb{C}_A^{\text{jit}}\}$, and implements the recursive nature of Crash types. We include the connectives \vee and \rightsquigarrow solely for the purpose of defining Crash types; they are not used elsewhere. Defining Crash types using these connectives will allow us to define the logical relation in Sec. 5 based on the intended meaning of its index type. Some well-formed types, e.g., $\text{nat} \rightsquigarrow \text{nat} \rightsquigarrow \uparrow \text{unit}$, are not accepted by our type system introduced in Sec. 4.4. These types have no inhabitants, i.e., no well-typed configuration is of these types.

$$\begin{array}{ll}
 \text{store types } A := \text{int} \mid \text{bool} & \text{stable types } \tau^s := \text{nat} \rightsquigarrow \tau^s \mid \uparrow \tau^u \\
 \text{basic types } T := \text{unit} \mid A & \text{unstable types } \tau^u := T \mid \downarrow \tau^s \mid \tau^u \vee \tau^u \mid v_t \\
 \text{Volatile store typing context } \Sigma := \cdot \mid x : \downarrow_u^s \uparrow_u^s A @ Ck, \Sigma & \\
 \text{Nonvolatile store typing context } \Omega := \cdot \mid x : \uparrow_u^s A @ Rd, \Omega \mid x_{\text{ck}} : \uparrow_u^s A @ CK, \Omega & \\
 & \mid x : \uparrow_u^s A @ CK, \Omega
 \end{array}$$

A nonvolatile store typing context Ω assigns stable types to nonvolatile location variables, i.e. all variables in Ω have a type of the form $\uparrow_u^s A$. A volatile store typing context Σ assigns unstable types to volatile location variables, i.e., variables in Σ are of the type $\downarrow_u^s \uparrow_u^s A$. x_{ck} refers to a location that has been checkpointed. In the atomic mode, x_{ck} has an active volatile log in Σ .

Typing Judgments. Table 1 summarizes all the typing judgments. These judgments are parameterized over the execution mode `Md` of the expression or command to be typed. The judgment also tracks a variable b corresponding to the current energy level of this execution. b ranges over natural numbers (`nat`) and is constrained by a relation $\mathcal{R} \in \{\geq, >\}$ or is set to 0; where $b \geq 0$ is unconstrained. The constraint on b determines whether or not a command can evaluate a value without power failure. There are three judgments for command typing. The first judgment is used when the command has not yet successfully finished

$$\begin{array}{c}
\text{jit} \mid b \geq 0 : \text{nat} \mid \Omega; \cdot \vdash_{\emptyset} c : \mathbf{C}_{\text{unit}} \quad b : \text{nat} \mid \Omega \vdash p : \uparrow \mathbf{C}_{\text{unit}} \\
\hline
b : \text{nat} \mid \Omega \vdash c; p : \uparrow \mathbf{C}_{\text{unit}} \quad (\text{T-P-SEQ}) \\
\\
\Omega_0 \mid \Sigma_0 = \text{InitWorld}_t(\Omega; \rho) \\
\text{Sig} = \{\text{alD}(c_0) \mid b \geq 0 : \text{nat} \mid \Omega_0; \Sigma_0 \vdash c_0 : \mathbf{C}_{\text{unit}}\} \\
\text{alD}(c_0) \mid b \geq 0 : \text{nat} \mid \Omega_0; \Sigma_0 \vdash_{\text{Sig}} c_0 : \mathbf{C}_{\text{unit}} \quad b : \text{nat} \mid \Omega \vdash p : \uparrow \mathbf{C}_{\text{unit}} \\
\hline
b : \text{nat} \mid \Omega \vdash \text{Ckpt}[\text{alD}, \rho](c_0); p : \uparrow \mathbf{C}_{\text{unit}} \quad (\text{T-P-CKPT})
\end{array}$$

Fig. 7. Program typing

executing; its next step, depending on its constraint \mathcal{R} , may or may not crash. When the command reaches type $\downarrow \uparrow \text{unit}$, b no longer needs to be constrained as the execution succeeded without power failure. The second judgment invokes the third judgment to type the configuration after the volatile log is committed: in the typing rule for committing the volatile log, the conclusion is of the form of the second judgment and the premise is of the form of the third. For expression typing, we distinguish expressions on the right of an assignment (being read) from those on the left of an assignment (being written to) via subscripts RD and WT, respectively. The expressions that are being written to are only of the simple form x . As no execution is required to evaluate x , we consider its judgment crash free, so no constraint is required on b . For program typing, we only have one judgment that refers to the type of the program before the execution of its next block starts. The rest of the judgments type states after a crash. The first judgment uses the constraint $b = 0$, which corresponds to the power failure condition. It invokes the second judgment, which types a state right after crash. The third judgment types the state awaiting energy to continue re-execution, and the final judgment types the state that is ready for restoration and re-execution.

4.4 Typing Rules

Program Typing. Fig. 7 shows the typing rules for programs. The P-SEQ rule types program $c; p$ by first typing c under jit mode, requiring $b \geq 0$, and then typing the rest of the program. The volatile memory context is empty for now, but will be populated when the let commands allocate new volatile locations.

The P-CKPT rule types the command c_0 enclosed in an atomic region under the mode $\text{alD}(c_0)$ and then types the rest of the program p . The first premise sets up the initial typing contexts for nonvolatile and volatile memories, as illustrated in Fig. 2. The partial function InitWorld_t initializes the volatile memory by creating a log of variables in Ω that are not read-only. Ω can be uniquely split into Ω^c and Ω^r , where Ω^r is the set of all read-only locations in Ω , and Ω^c is the set of all locations that are not read-only. This function is defined below:

$$\begin{array}{l}
\Omega_0 \mid \Sigma_0 = \text{InitWorld}_t(\Omega; \rho) \text{ iff } \rho \subseteq \text{dom}(\Omega), \Omega_0 = \Omega^r, \Omega_{\text{ck}}^c \text{ and } \Sigma_0 = \downarrow \Omega^c \\
\text{where } \Omega = \Omega^c, \Omega^r \text{ and } \Omega^r = \Omega \setminus \rho.
\end{array}$$

Here $\Omega^r = \Omega \setminus \rho$ is a subset of Ω where locations are declared in ρ to be read-only, and Ω^c are all other locations in Ω . The context Ω_{ck}^c , is defined as

$\Omega_{\text{ck}}^c = \{x_{\text{ck}} : \uparrow A @ q \mid x : \uparrow A @ q \in \Omega^c\}$, and the context $\downarrow \Omega^c$, is defined as $\downarrow \Omega^c = \{x : \downarrow \uparrow A @ q \mid x : \uparrow A @ q \in \Omega^c\}$. If the set of read only variables, ρ , is not in the domain of Ω , then the function InitWorld_t is not defined.

In rules P-SEQ and P-CKPT, the command typing judgment in the premise makes use of a signature (subscripts \emptyset and Sig , respectively) to type check the command relative to the signature. The signature is populated at different stages of type checking the JIT and atomic regions. In an atomic region, rule T-P-CKPT populates the signature at the beginning of the region with the initial judgment which includes the region's original command c_0 and static memory context $\Omega_0; \Sigma_0$. The region is then typed relative to the signature. In JIT mode, the signature is populated later with the judgment just at the point of the failure (rule T-ENOUGH?). The program remembers that it built a typing derivation for the judgment in the signature such that when it restores from a power failure, it refers to the signature and checks that the restored judgment matches the one stored in the signature without needing to derive it again. This makes the typing derivations finitary and inductive.

Command and Expression Typing. Fig. 8 shows selected typing rules for commands. The T-SKIP rule declares the command `skip` as the stable type $\uparrow \text{unit}$. Rule T-V-SUCC applies when the command successfully completes its execution and still has one unit of energy available ($b > 0$) to conclude the execution. In this case, we close off the energy level variable and continue typing the command against the type $\downarrow \uparrow \text{unit}$. Rule T-C-SHIFT is invoked by T-V-SUCC and updates the memory typing contexts by removing checkpointed locations in Ω as now they are not needed, and making locations in Σ stable as now they are committed. This corresponds to the last step of Fig. 2.

The rules T-LET and T-ASSIGN, are mostly standard except that we consider crashes. For example, in typing the assign command $x := e$, the first premise of T-ASSIGN considers the type of expression e to be the Crash type C_A^{Md} , but in the second premise we require the location x to be of type $\downarrow \uparrow A$, i.e., the location only considers the type corresponding to the case where execution of e can be completed successfully. The reason is that the assignment only occurs if the execution of e is successful. The constraint on the energy levels for premises goes back to $b \geq 0$, as we use one energy unit to deconstruct these commands.

The rule T-ENOUGH? checks two premises based on the value of $b \geq 0$. The third premise, a crash judgment, corresponds to the case where $b = 0$ (typing rules for crash judgments are given later in this section) and the fourth premise corresponds to the case where $b > 0$. The condition $b > 0$ states that there is at least one unit of energy available to decompose one command construct, e.g., via T-LET or T-ASSIGN. This rule populates the signature for JIT commands. The second premise states that the signature remains intact if the mode is atomic, but is populated by Sig' if the mode is JIT. In the JIT mode, after a power failure, the command c is restored to itself, and Sig' remembers that the well-typedness of the command when the energy level is non-negative has been checked already.

Expression typing rules are very similar to those of the commands. Fig. 8 shows a few selected rules. The T-LOC-WRITE and T-LOC-READ rules match

Commands

$$\begin{array}{c}
\frac{}{\text{Md} \mid b : \text{nat} \mid \Omega \vdash_{\text{sig}} \text{skip} : \uparrow \text{unit}} \text{(T-SKIP)} \\
\\
\frac{\Sigma = \downarrow \Sigma' \quad \Omega = \Omega', \Omega'_{\text{ck}} \quad \text{Md} \mid b : \text{nat} \mid \Omega', \Sigma' \vdash_{\text{sig}} \text{skip} : \uparrow \text{unit}}{\text{Md} \mid b : \text{nat} \mid \Omega; \Sigma \vdash_{\text{sig}} \text{skip} : \downarrow \uparrow \text{unit}} \text{(T-C-SHIFT)} \\
\\
\frac{\text{Md} \mid b : \text{nat} \mid \Omega; \Sigma \vdash_{\text{sig}} \text{skip} : \downarrow \uparrow \text{unit}}{\text{Md} \mid b > 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{sig}} \text{skip} : \tau \vee \downarrow \uparrow \text{unit}} \text{(T-V-SUCC)} \\
\\
\frac{\text{Md} \mid b \geq 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{RD}; \text{sig}} e_1 : \mathbf{C}_A^{\text{Md}} \quad \text{Md} \mid b \geq 0 : \text{nat} \mid \Omega; \Sigma, x : \downarrow \uparrow A @ \text{CK} \vdash_{\text{sig}} c : \tau}{\text{Md} \mid b > 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{sig}} \text{let } x = e_1 \text{ in } c : \tau} \text{(T-LET)} \\
\\
\frac{\text{Md} \mid b \geq 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{RD}; \text{sig}} e : \mathbf{C}_A^{\text{Md}} \quad \text{Md} \mid b > 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{WT}} x : \downarrow \uparrow A}{\text{Md} \mid b > 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{sig}} x := e : \mathbf{C}_{\text{unit}}^{\text{Md}}} \text{(T-ASSIGN)} \\
\\
\frac{\text{Md} \mid b = 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{sig}''} c : \tau \quad \text{Md} \mid b > 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{sig}} c : \tau}{\text{Md} \mid b \geq 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{sig}} c : \tau} \text{(T-ENOUGH?) } \\
\text{Sig}' = \{\text{Md} \mid b \geq 0 : \text{nat} \mid \Omega; \Sigma \vdash c : \tau\} \\
\text{Sig}'' = \text{if Md} = \text{jit, then Sig}', \text{ else Sig}
\end{array}$$

Expressions

$$\begin{array}{c}
\frac{\Omega, \Sigma' = x : \uparrow A @ q, \Omega'_2 \quad q \neq \text{RD}}{\text{Md} \mid b : \text{nat} \mid \Omega, \Sigma' \vdash_{\text{Wt}} x : \uparrow A} \text{(T-LOC-WRITE)} \\
\\
\frac{\Omega = x : \uparrow A @ q, \Omega'}{\text{Md} \mid b : \text{nat} \mid \Omega \vdash_{\text{RD}} x : \uparrow A} \text{(T-LOC-READ)} \qquad \frac{}{\text{Md} \mid b : \text{nat} \mid \Omega \vdash_{\text{RD}} \text{tt} : \uparrow \text{bool}} \text{(T-BOOL-T)}
\end{array}$$

Fig. 8. Selected command and expression typing

the location variable x with an existing variable inside the context. T-LOC-WRITE performs an extra check to make sure that x is not a read-only variable.

Statement typing Fig. 9 presents the typing rules for crash instructions. The crash is detected by the depleted energy level $b = 0$ in the T-V-CRASH rule. In the premise, the crash instruction $\downarrow \varepsilon \# \text{in}(b > 0, \uparrow \kappa')$ is typed. In JIT mode, the T-JIT-STOP rule brings a checkpointed version of all the volatile variables in Σ inside Ω since they are checkpointed then. In atomic mode, T-AID-STOP rule simply drops the volatile locations in Σ . The T-CHARGE rule inputs a new energy level from the energy channel ε , regardless of the mode. The first premise shows that the energy channel is needed to provide a natural number greater than zero. Finally, the T-JIT-RESTORE and T-AID-RESTORE rules prepare and check rebooted system in JIT and atomic modes, respectively. In both modes, volatile memory is restored from the checkpointed locations in Ω . In the atomic mode, the checkpointed locations persist in Ω as we may need them for the

$$\begin{array}{c}
 \frac{\text{Md} \mid \cdot \mid \Omega; \Sigma \vdash_{\text{Sig}} \downarrow \varepsilon \# \text{in}(b > 0, \uparrow \kappa') : \downarrow (\text{nat} \rightsquigarrow \uparrow \mathbf{C}_T^{\text{Md}})}{\text{Md} \mid b = 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{Sig}} \kappa' : \downarrow (\text{nat} \rightsquigarrow \uparrow \mathbf{C}_T^{\text{Md}}) \vee \downarrow \uparrow T} \quad (\text{T-V-CRASH}) \\
 \\
 \frac{\Sigma = \downarrow \uparrow \Sigma' \quad \text{jit} \mid \cdot \mid \Omega, \uparrow \Sigma'_{\text{ck}} \vdash_{\text{Sig}} \varepsilon \# \text{in}(b > 0, \uparrow \kappa') : (\text{nat} \rightsquigarrow \uparrow \mathbf{C}_T^{\text{s}})}{\text{jit} \mid \cdot \mid \Omega; \Sigma \vdash_{\text{Sig}} \downarrow \varepsilon \# \text{in}(b > 0, \uparrow \kappa') : \downarrow (\text{nat} \rightsquigarrow \uparrow \mathbf{C}_T^{\text{s}})} \quad (\text{T-JIT-STOP}) \\
 \\
 \frac{\text{aID}(c_0) \mid \cdot \mid \Omega \vdash_{\text{Sig}} \varepsilon \# \text{in}(b > 0, \uparrow \kappa') : (\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}^{\text{s}})}{\text{aID}(c_0) \mid \cdot \mid \Omega; \Sigma \vdash_{\text{Sig}} \downarrow \varepsilon \# \text{in}(b > 0, \uparrow \kappa') : \downarrow (\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}^{\text{s}})} \quad (\text{T-AID-STOP}) \\
 \\
 \frac{\varepsilon \# \text{in}() : \text{nat} > 0 \quad \text{Md} \mid b > 0 : \text{nat} \mid \Omega \vdash_{\text{Sig}} \uparrow \kappa' : \uparrow \mathbf{C}_T^{\text{s}}}{\text{Md} \mid \cdot \mid \Omega \vdash_{\text{Sig}} \varepsilon \# \text{in}(b > 0, \uparrow \kappa') : (\text{nat} \rightsquigarrow \uparrow \mathbf{C}_T^{\text{s}})} \quad (\text{T-CHARGE}) \\
 \\
 \frac{\Omega = \Omega', \Omega''_{\text{ck}} \quad \text{jit} \mid b \geq 0 : \text{nat} \mid \Omega'; \downarrow \Omega'' \vdash \kappa' : \mathbf{C}_T \in \mathbf{Sig}}{\text{jit} \mid b > 0 : \text{nat} \mid \Omega \vdash_{\text{Sig}} \uparrow \kappa' : \uparrow \mathbf{C}_T} \quad (\text{T-JIT-RESTORE}) \\
 \\
 \frac{\Omega = \Omega', \Omega''_{\text{ck}} \quad \text{aID}(c_0) \mid b \geq 0 : \text{nat} \mid \Omega; \downarrow \Omega'' \vdash c_0 : \mathbf{C}_{\text{unit}} \in \mathbf{Sig}}{\text{aID}(c_0) \mid b > 0 : \text{nat} \mid \Omega \vdash_{\text{Sig}} \uparrow \kappa' : \uparrow \mathbf{C}_{\text{unit}}} \quad (\text{T-AID-RESTORE})
 \end{array}$$

Fig. 9. Crash, restore, and checkpoint typing

next power failure. Alternatively, in the JIT mode, checkpoints are dropped from Ω and execution continues with the expression or command κ , which was running right before the crash. In the atomic mode, execution continues with the original command c_0 enclosed in the atomic region. Instead of retyping the restored judgments, we check if there are already typing derivations by matching them up with the saved judgment in the signature.

5 Logical Relation for Intermittent Execution

We establish a logical relation to prove idempotency, which states that every intermittent execution of a program can be simulated by a continuous execution. The logical relation relates an intermittent execution with a continuous one and is indexed by Crash types. A continuous run is one with an infinite energy level, ∞ . Crash types are recursive, yielding possible infinite atomic region re-executions. Thus, we use the maximum number of executions (also power failures) as a step index to stratify our logical relation to ensure its well-foundedness.

The logical relation (defined in Sec. 5.1) relies on **PwOff**, **Restore**, and **Commit** functions, referred to as power failure, restore, and commit policies, respectively. We establish specific policies for atomic and JIT execution modes. We formalize *semantic typing* as every atomic and JIT region of the program being logically-related to themselves. We prove that the semantically well-typed programs are idempotent across power failures in Sec. 5.2. The definitions match the memory operations in the dynamic rules that deal with crash, restore, and re-execution (D-S-AID/ D-S-JIT, D-R-AID/ D-R-JIT, and D-P-CKPT/

$\text{Md} \mid b \geq 0 : \text{nat} \mid \Omega \mid \Sigma \Vdash c_1 \leq c_2 : \mathbf{C}_{\text{unit}}$
 iff $\forall n, m \geq 0. \forall \gamma, \text{NV}, \mathbf{V}. \text{s.t. } \text{NV} \mid \mathbf{V} \Vdash \gamma :: \Omega \mid \Sigma.$
 $(\gamma \mid \text{Md} \mid n \mid \text{NV} \mid \mathbf{V} \mid c_1, \gamma \mid \text{Md} \mid \infty \mid \text{NV} \mid \mathbf{V} \mid c_2) \in \mathcal{E}[\mathbf{C}_{\text{unit}}]^m$

Term Relation

$\mathcal{E}[\mathbf{C}_{\text{unit}}]^{m+1} = \{(\gamma_1 \mid \text{Md} \mid n_1 \mid \text{NV}_1 \mid \mathbf{V}_1 \mid c_1, \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2) \text{ s.t.}$
 $\exists. (\gamma'_1 \mid \text{Md}' \mid n'_1 \mid \text{NV}'_1 \mid \mathbf{V}'_1 \mid c'_1) \text{ s.t.}$
 $\gamma_1 \mid \text{Md} \mid n_1 \mid \text{NV}_1 \mid \mathbf{V}_1 \mid c_1 \rightarrow_{\text{irred}}^* \gamma'_1 \mid \text{Md}' \mid n'_1 \mid \text{NV}'_1 \mid \mathbf{V}'_1 \mid c'_1 \wedge$
 $\exists. (\gamma'_2 \mid \text{Md}' \mid \infty \mid \text{NV}'_2 \mid \mathbf{V}'_2 \mid c'_2) \text{ s.t.}$
 $\gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2 \rightarrow^* \gamma'_2 \mid \text{Md}' \mid \infty \mid \text{NV}'_2 \mid \mathbf{V}'_2 \mid c'_2 \wedge$
 $(\gamma'_1 \mid \text{Md}' \mid n'_1 \mid \text{NV}'_1 \mid \mathbf{V}'_1 \mid c'_1, \gamma'_2 \mid \text{Md}' \mid \infty \mid \text{NV}'_2 \mid \mathbf{V}'_2 \mid c'_2) \in \mathcal{V}[\mathbf{C}_{\text{unit}}]^{m+1}\}$
 $\mathcal{E}[\mathbf{C}_{\text{unit}}]^0 = \{(\gamma_1 \mid \text{Md} \mid n_1 \mid \text{NV}_1 \mid \mathbf{V}_1 \mid c_1, \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2)\}$

Value Relation

$\mathcal{V}[\uparrow \mathbf{unit}]^m = \{(\gamma \mid \text{Md} \mid n_1 \mid \text{NV}_1 \mid \text{skip}, \gamma \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \text{skip}) \text{ s.t. } \text{NV}_1 = \text{NV}_2\}$
 $\mathcal{V}[\downarrow \uparrow \mathbf{unit}]^m = \{(\gamma_1 \mid \text{Md} \mid n_1 \mid \text{NV}_1 \mid \mathbf{V}_1 \mid \text{skip}, \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid \text{skip}) \text{ s.t.}$
 $\text{Commit}(\gamma_i \mid \text{Md} \mid \text{NV}_i \mid \mathbf{V}_i) = \gamma'_i \mid \text{NV}'_i \wedge$
 $(\gamma'_1 \mid \text{Md} \mid n_1 \mid \text{NV}'_1 \mid \text{skip}, \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}'_2 \mid \text{skip}) \in \mathcal{V}[\uparrow \mathbf{unit}]^m\}$
 $\mathcal{V}[\uparrow \mathbf{C}_{\text{unit}}]^m = \{(\gamma_1 \mid \text{Md} \mid n \mid \text{NV}_1 \mid \uparrow \kappa, \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2) \text{ s.t.}$
 $\text{restore}(\gamma_1, \text{Md}, \text{NV}_1, \kappa) = \text{NV}_0 \mid \mathbf{V}_0 \mid c_0 \wedge$
 $(\gamma_1 \mid \text{Md} \mid n \mid \text{NV}_0 \mid \mathbf{V}_0 \mid c_0, \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2) \in \mathcal{E}[\mathbf{C}_{\text{unit}}]^m\}$
 $\mathcal{V}[\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}]^m = \{(\gamma_1 \mid \text{Md} \mid \cdot \mid \text{NV}_1 \mid \varepsilon \# \text{in}(n > 0, \uparrow \kappa), \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2) \text{ s.t.}$
 $\forall n > 0. (\gamma_1 \mid \text{Md} \mid n \mid \text{NV}_1 \mid \uparrow \kappa, \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2) \in \mathcal{V}[\uparrow \mathbf{C}_{\text{unit}}]^m\}$
 $\mathcal{V}[\downarrow (\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}})]^m = \{(\gamma_1 \mid \text{Md} \mid \cdot \mid \text{NV}_1 \mid \mathbf{V}_1 \mid \downarrow \varepsilon \# \text{in}(n > 0, \uparrow \kappa), \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2)$
 $\text{s.t. } \text{PwOff}(\gamma_1, \text{Md}, \text{NV}_1, \mathbf{V}_1) = \gamma'_1 \mid \mathbf{V}' \wedge$
 $(\gamma'_1 \mid \text{Md} \mid \cdot \mid \mathbf{V}', \text{NV}_1 \mid \varepsilon \# \text{in}(n > 0, \uparrow \kappa), \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2)$
 $\in \mathcal{V}[\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}]^m\}$
 $\mathcal{V}[\mathbf{C}_{\text{unit}}]^{m+1} = \{(\gamma_1 \mid \text{Md} \mid n_1 \mid \text{NV}_1 \mid \mathbf{V}_1 \mid c_1, \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2)$
 s.t. either
 $n_1 = 0 \wedge (\gamma_1 \mid \text{Md} \mid \cdot \mid \text{NV}_1 \mid \mathbf{V}_1 \mid \downarrow \varepsilon \# \text{in}(n_1 > 0, \uparrow \kappa),$
 $\gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2) \in \mathcal{V}[\downarrow (\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}})]^m, \text{ or}$
 $n_1 > 0 \wedge (\gamma_1 \mid \text{Md} \mid n_1 \mid \text{NV}_1 \mid \mathbf{V}_1 \mid c_1, \gamma_2 \mid \text{Md} \mid \infty \mid \text{NV}_2 \mid \mathbf{V}_2 \mid c_2)$
 $\in \mathcal{V}[\downarrow \uparrow \mathbf{unit}]^m\}$

Fig. 10. Logical relation

D-P-SEQ) for atomic and JIT regions, We prove that our syntactically well-typed programs are semantically well-typed. We generalize semantic typing rules, allowing custom power failure, restore, and commit policies (Sec. 5.3).

5.1 Semantic Typing via a Logical Relation

The logical relation, written $\text{Md} \mid b \geq 0 : \text{nat} \mid \Omega \mid \Sigma \Vdash c_1 \leq c_2 : \mathbf{C}_{\text{unit}}$, is defined in Fig. 10 by a lexicographic induction on the index m and the structure of the

types. The judgment $\text{NV} \mid \mathbb{V} \Vdash \gamma :: \Omega \mid \Sigma$ in the definition states that γ maps the variables in Σ and Ω to locations in \mathbb{V} and NV resp., such that their qualifiers and types match. Similar to prior work [2,16,42], our definition consists of a term relation $\mathcal{E}[\llbracket \mathbf{C}_{\text{unit}} \rrbracket]^m$ and a value relation $\mathcal{V}[\llbracket \tau \rrbracket]^m$.

Term Relation. A pair of open command configurations of type \mathbf{C}_{unit} are in the term relation of index m if any intermittent execution of the first one after m power failures is indistinguishable from a continuous execution of the second one. In particular, for index $m+1$, the term relation relates two configurations at type \mathbf{C}_{unit} if the first configuration eventually steps to a value (or “irreducible”) configuration, i.e., it either evaluates to `skip` or its energy level depletes ($n'_1 = 0$), and the second configuration can take zero or more steps such that the pair continue to be in the value relation of $\mathcal{V}[\llbracket \mathbf{C}_{\text{unit}} \rrbracket]^{m+1}$. When the index is $m = 0$, no execution is observed, so any two configurations are in the term relation. Here, *irred* refers to $\gamma'_1 \mid \mathbf{Md}' \mid n'_1 \mid \text{NV}'_1 \mid \mathbb{V}'_1 \mid c'_1$ being an irreducible configuration, i.e. it cannot take any more steps. Since our semantics for commands is deterministic, for each configuration $\gamma_1 \mid \mathbf{Md} \mid n_1 \mid \text{NV}_1 \mid \mathbb{V}_1 \mid c_1$ there is exactly one such irreducible configuration.

Value Relation. The value relation is defined based on the intended meaning of the type, and relates two value configurations that will have the same effect on the stores. The value relation relates two open command configurations at type \mathbf{C}_{unit} and index $m+1$ if either (a) the first configuration has faced a power failure, and the two configurations continue to relate by $\mathcal{V}[\llbracket \downarrow(\mathbf{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}) \rrbracket]^m$, or (b) the first configuration executed successfully without any power failures, and the two configurations are related by $\mathcal{V}[\llbracket \downarrow \uparrow \mathbf{unit} \rrbracket]^m$. This definition matches the disjunctive nature of type \mathbf{C}_{unit} , which is recursively defined in the signature as $\downarrow(\mathbf{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}) \vee \downarrow \uparrow \mathbf{unit}$. Since we unfold the recursive definition of \mathbf{C}_{unit} , we decrease the index from $m+1$ to m to ensure the relation’s well-foundedness. Note that the value relation is neither defined nor called for \mathbf{C}_{unit} at index 0.

The value relations in the third, fourth, and fifth rows of Fig. 10 are defined based on the type of the *first configuration*; the second configurations in these relations continue to be of type \mathbf{C}_{unit} . Only in the relations defined in the first and second rows of Fig. 10 do the types of both configurations match the indexed type of the relation. Hence, the value relation has varying arity: in the first and second rows of Fig. 10, the relation is *binary* while in the rest, the relation degenerates to *unary*, with the second configuration as its Kripke world [18].

The value relation at type $\downarrow(\mathbf{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}})$ relates two configurations if the first one runs the crash instruction $\downarrow \varepsilon \# \text{in}(n > 0, \uparrow \kappa)$ and a power failure policy creates a checkpoint of volatile locations such that the configurations continue to be in the value relation at type $(\mathbf{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}})$. The power failure function in an atomic mode is defined to checkpoint none of the volatile locations, i.e., $\text{PwOfff}(\gamma, \text{aID}(c_0), \text{NV}_1, \mathbb{V}_1) = \gamma' \mid \emptyset$, where γ' is the largest restriction of γ with $\text{range}(\gamma') = \text{dom}(\text{NV}_1)$, and defined to checkpoint all volatile locations in JIT mode, i.e., $\text{PwOfff}(\gamma, \text{jit}, \text{NV}_1, \mathbb{V}_1) = \gamma \mid \mathbb{V}_1$.

The value relation at type $(\mathbf{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}})$ is defined similarly to a function type in a value relation and requires the configurations to be related at type $(\uparrow \mathbf{C}_{\text{unit}})$ for every energy input level n provided to the first configuration.

The value relation at type $\uparrow \mathbf{C}_{\text{unit}}$ requires the first configuration to run the crash instruction $\uparrow \kappa$. The defined restore policy restores the nonvolatile memory \mathbf{NV}_0 , volatile memory \mathbf{V}_0 , and re-execution command c_0 such that the configurations continue to be related in the term interpretation at type \mathbf{C}_{unit} . In an atomic mode, the restore function is defined as $\mathbf{restore}(\gamma, \mathbf{aID}(c), \mathbf{NV}_1, \kappa) = \mathbf{NV}_1 \mid \mathbf{NV}'' \mid c$ where $\mathbf{NV}_1 = \mathbf{NV}', \mathbf{NV}''_{\text{ck}}$. In the JIT mode, the restore function is defined as $\mathbf{restore}(\gamma, \mathbf{jit}(c), \mathbf{NV}_1, \kappa) = \mathbf{NV}'_1 \mid \mathbf{NV}'' \mid c$ where $\mathbf{NV}_1 = \mathbf{NV}', \mathbf{NV}''_{\text{ck}}$. We write $\mathbf{NV}_1 = \mathbf{NV}', \mathbf{NV}''_{\text{ck}}$ to state that \mathbf{NV}_1 can be uniquely partitioned into all locations $(\mathbf{NV}''_{\text{ck}})$ that are checkpointed, i.e., of the form ℓ_{ck} , and regular locations (\mathbf{NV}') of the form ℓ . \mathbf{NV}'' is the non-checkpointed version of $\mathbf{NV}''_{\text{ck}}$ which could be retrieved by removing the ck subscript from every location in $\mathbf{NV}''_{\text{ck}}$.

The value relation at type $\downarrow \uparrow \mathbf{unit}$ requires both configurations to run \mathbf{skip} , and the defined commit policy creates nonvolatile memories for both runs such that they continue to be related at type $\uparrow \mathbf{unit}$. In an atomic mode, the commit function is defined to replace the checkpointed locations in the nonvolatile memory with their volatile log, i.e., $\mathbf{Commit}(\gamma \mid \mathbf{aID}(c_0) \mid \mathbf{NV}_1 \mid \mathbf{V}_1) = \gamma' \mid \mathbf{NV}'_1 \mid \mathbf{V}''$, where $\mathbf{NV}_1 = \mathbf{NV}'_1, \mathbf{NV}''_{\text{ck}}$ and $\mathbf{V}_1 = \mathbf{V}'_1, \mathbf{V}''$ and $\text{dom}(\mathbf{V}'') = \text{dom}(\mathbf{NV}'')$. Moreover, $\gamma' \subseteq \gamma$, with $\text{range}(\gamma') = \text{dom}(\mathbf{NV}_1) \cup \text{dom}(\mathbf{V}'')$. In the JIT mode, the commit function simply drops all volatile memory, i.e., $\mathbf{Commit}(\gamma \mid \mathbf{jit} \mid \mathbf{NV}_1 \mid \mathbf{V}_1) = \gamma' \mid \mathbf{NV}_1$, $\gamma' \subseteq \gamma$, with $\text{range}(\gamma') = \text{dom}(\mathbf{NV}_1)$.

The value relation at type $\uparrow \mathbf{unit}$ requires the successful executions to store the same values in their memories, i.e., $\mathbf{NV}_1 = \mathbf{NV}_2$.

Semantic Typing. A program is semantically well-typed if every JIT and atomic region of it is self-related under our logical relation.

$$\frac{\mathbf{jit} \mid b \geq 0 : \mathbf{nat} \mid \Omega; \cdot \Vdash c \leq c : \mathbf{C}_{\text{unit}} \quad b : \mathbf{nat} \mid \Omega \Vdash p : \uparrow \mathbf{C}_{\text{unit}}}{b : \mathbf{nat} \mid \Omega \Vdash c; p : \uparrow \mathbf{C}_{\text{unit}}} \text{ (P-SEQ-SEMANTIC)}$$

$$\frac{\Omega_0 \mid \Sigma_0 = \text{InitWorld}_t(\Omega; \rho) \quad \mathbf{aID}(c_0) \mid b \geq 0 : \mathbf{nat} \mid \Omega_0; \Sigma_0 \Vdash c_0 \leq c_0 : \mathbf{C}_{\text{unit}} \quad b : \mathbf{nat} \mid \Omega \Vdash p : \uparrow \mathbf{C}_{\text{unit}}}{b : \mathbf{nat} \mid \Omega \Vdash \text{Ckpt}[\mathbf{aID}, \rho](c_0); p : \uparrow \mathbf{C}_{\text{unit}}} \text{ (P-CKPT-SEMANTIC)}$$

5.2 Semantic Typing for Idempotency

The fundamental theorem of our logical relation states that syntactically well-typed programs are also semantically well-typed by proving that syntactically well-typed JIT and atomic regions are self-related. We state and prove the theorem in Sec. 6 but devote this section to explaining why being self-related implies idempotency. We explain it separately for JIT and atomic blocks.

Stepping a JIT block. Consider a program of form $[\chi_1 \triangleright \varepsilon] \otimes \gamma_1 \mid n \mid \mathbf{NV}_1 \mid c_1; p$ that can take a step to $[\chi_k \triangleright \varepsilon] \otimes \gamma \mid n'_k \mid \mathbf{NV}'_k \mid p$ via the D-P-SEQ rule. By the D-P-SEQ rule, we know that the command c_1 is successfully executed to completion with possibly m -many power failures along the way: $[\chi_1 \triangleright \varepsilon] \otimes \gamma_1 \mid$

$\text{jit} \mid n \mid \text{NV}_1 \mid \cdot \mid c_1 \Rightarrow^* [\chi_k \triangleright \varepsilon] \otimes \gamma'_k \mid \text{jit} \mid n'_k \mid \text{NV}'_k \mid \text{V}'_k \mid \text{skip}$. Our goal is to simulate this execution in a continuous setting. To model a continuous run, we run the configuration with ∞ , an energy level: $[\chi \triangleright \varepsilon] \otimes \gamma_1 \mid \text{jit} \mid \infty \mid \text{NV}_1 \mid \cdot \mid c_1 \Rightarrow^* [\chi \triangleright \varepsilon] \otimes \gamma'_j \mid \text{jit} \mid \infty \mid \text{NV}'_j \mid \text{V}'_j \mid \text{skip}$.

Fig. 11 shows the construction of the simulation. We start with the assumption that the configuration with n energy level is self-related when given energy level ∞ for every index, including $m + 1$ (point (1) in Fig. 11). We show that if the first configuration takes one or more steps, the second configuration can take zero or more steps so that the intermediate regions continue to relate.

By definition of the term interpretation, c_1 in the first configuration is executed until the first power failure occurs. Moreover, by the relation, we can execute c_1 in the second configuration, too, such that the resulting configurations remain related (point (2) in Fig. 11) by the value interpretation at type \mathbf{C}_{unit} . The first configuration takes a step from point (2) to point (3) using the D-CRASH rule by the computational semantics. By the definition of the logical relation, the two configurations continue to be related by the value interpretation at type $\downarrow(\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}})$. Then the first configuration takes a step from point (3) to point (4) by the D-S-JIT rule; in this case, we know (by the assumptions of the rule) $\text{V}' = \text{V}'_1$ and $\gamma'_1 = \gamma$. This matches the definition of the power-off policy for JIT blocks (see Sec. 5.1), and thus the two configurations remain related by the value relation at type $\text{nat} \rightsquigarrow \uparrow \mathbf{C}_{\text{unit}}$. Next, the first configuration takes a step to point (5) by inputting a new energy level from the environment (n_2). By the definition of the value relations, the two configurations will remain related by the value interpretation at type $\uparrow \mathbf{C}_{\text{unit}}$.

Finally, the configuration steps to point (6) by D-RESTORE-JIT that copies all checkpointed locations inside the volatile memory and continues by running the interrupted command κ , i.e., here $\text{NV}_0 = \text{NV}'_1$ and $\text{V}_0 = \text{V}' = \text{V}'_1$ and $c_0 = \kappa$. This matches the restore policy defined for JIT regions; thus, the configurations continue to be related by the *term relation* at type \mathbf{C}_{unit} , similar to what we had earlier at point (1) in Fig. 11, but with fewer power failures remaining.

Now, when the first configuration finally steps to point (8), by the definition of the logical relation, we know that the second configuration steps into skip too. Thus, we can apply the D-Ckpt rule on the second configuration. The volatile memory V'_j is dropped, and the mapping is reset to γ , i.e., it matches the commit policy defined for JIT blocks. in the logical relation. By Fig. 11-d, we get $\text{NV}'_j = \text{NV}'_k$, which completes deriving our goal.

Stepping an atomic region. We can build the desired simulation by taking the same steps described for a JIT region. Similarly, the key point is that the power-off and restore policies exactly match how the rules D-S-AID and D-RESTORE-AID, respectively, handle nonvolatile and volatile memories, and the commit policy corresponds to the `FinWorld` function in the D-CKPT rule.

We showed that our logical relation ensures idempotency for JIT and atomic regions. In the next section, we show that our logical relation formalizes a semantic typing to ensure idempotency of more general policies.

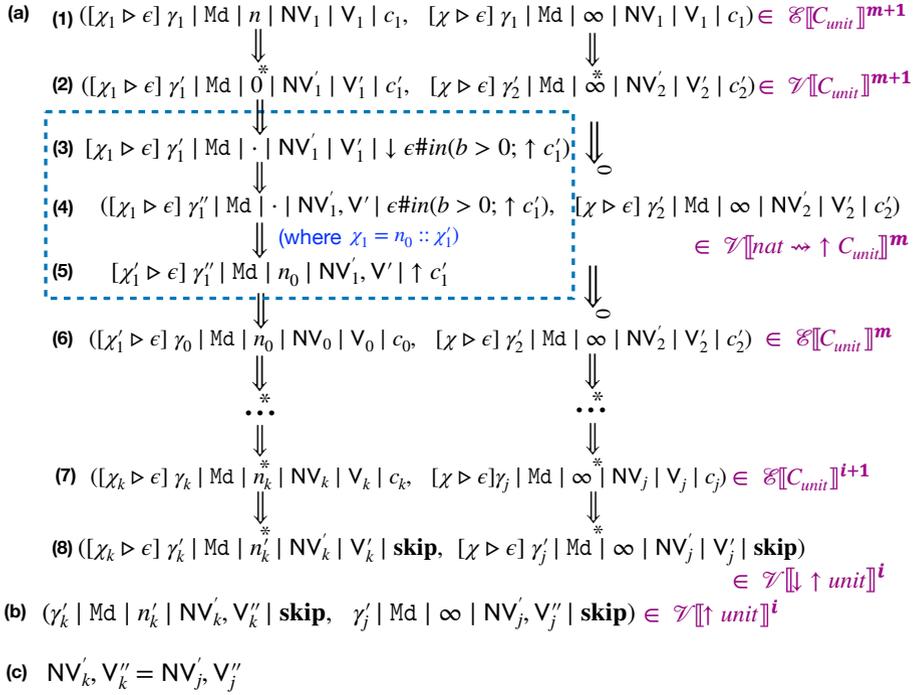


Fig. 11. Why the logical relation is enough.

5.3 More General Policies

We utilize our semantic typing to allow custom policies for power failure, restore, and commit. We extend the grammar of programs as $p := \cdot \mid \text{Reg}[\text{aID}, \overrightarrow{\text{arg}}](c); p$, where $\overrightarrow{\text{arg}}$ refers to the arguments that the programmer decides to pass to the region for initialization. To each region, we assign a unique identifier aID that is associated with the three policies and two functions InitGeneral_t and InitGeneral_d to initialize the static and dynamic memories, respectively. We add the following semantic typing rule for the general regions:

$$\frac{c_0 \mid \Omega_0 \mid \Sigma_0 = \text{InitGeneral}_t(\Omega; \text{aID}; c; \overrightarrow{\text{arg}}) \quad \text{aID}(c_0) \mid b \geq 0 : \mathbf{nat} \mid \Omega_0; \Sigma_0 \Vdash c_0 \leq c_0 : \mathbf{C}_{\text{unit}} \quad b : \mathbf{nat} \mid \Omega \Vdash p : \uparrow \mathbf{C}_{\text{unit}}}{b : \mathbf{nat} \mid \Omega \Vdash \text{Reg}[\text{aID}, \overrightarrow{\text{arg}}](c); p : \uparrow \mathbf{C}_{\text{unit}}} \quad (\text{P-REG-SEMANTIC})$$

For a self-related region to be idempotent, its policies **Commit**, **PwOff**, and **Restore** must match the dynamics, so we add dynamic rules for custom regions in Fig. 12. The JIT and atomic region policies and their dynamic rules are instances of these general policies. As an example, the programmer can customize the policies of the first block of Fig. 1 to not checkpoint variable u . The program remains idempotent as the atomic region never reads u before writing to it. This

$$\begin{array}{c}
 \frac{\gamma_0 \mid \mathbf{NV}_0 \mid \mathbf{V}_0 \mid c_0 = \text{restore}(\mathbf{NV}, \mathbf{V}, \kappa, \mathbf{Md}, \gamma)}{[\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \uparrow \kappa \Rightarrow [\chi \triangleright \varepsilon] \otimes \gamma_0 \mid \mathbf{Md} \mid n \mid \mathbf{NV}_0 \mid \mathbf{V}_0 \mid c_0} \text{ (D-R-REG)} \\
 \\
 \frac{\begin{array}{l} n > 0 \quad \text{InitGeneral}_d(\mathbf{NV}; \mathbf{aID}; c; \gamma; \overline{\text{ar}}\hat{g}) = c_0, \mathbf{NV}_0, \mathbf{V}_0 \\ [\chi \triangleright \varepsilon] \otimes \mathbf{aID}(c_0) \mid n \mid \mathbf{NV}_0 \mid \mathbf{V}_0 \mid c_0 \Rightarrow^* [\chi' \triangleright \varepsilon] \otimes \mathbf{aID}(c_0) \mid n' \mid \mathbf{NV}' \mid \mathbf{V}' \mid \text{skip} \\ n' > 0 \quad \mathbf{NV}_1 = \text{Commit}(\mathbf{NV}'; \mathbf{V}'; \mathbf{aID}; \overline{\text{ar}}\hat{g}) \end{array}}{[\chi \triangleright \varepsilon] \otimes \gamma \mid n \mid \mathbf{NV} \mid \text{Reg}[(\mathbf{aID}; \overline{\text{ar}}\hat{g})](c); p \Rightarrow [\chi' \triangleright \varepsilon] \otimes \gamma \mid n' \mid \mathbf{NV}_1 \mid p} \text{ (D-REG)} \\
 \\
 \frac{\mathbf{V}' = \text{PwOff}(\mathbf{NV}, \mathbf{V}, \mathbf{Md}, \gamma)}{[\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid \cdot \mid \mathbf{NV} \mid \mathbf{V} \mid \downarrow \varepsilon \# \text{in}(b > 0; \uparrow \kappa) \Rightarrow [\chi \triangleright \varepsilon] \otimes \gamma \mid \mathbf{Md} \mid \cdot \mid \mathbf{NV}, \mathbf{V}' \mid \varepsilon \# \text{in}(b > 0; \uparrow \kappa)} \text{ (D-S-REG)}
 \end{array}$$

Fig. 12. Custom dynamic rules

policy is implemented by real systems [23,24,41]. Our static typing rules can be extended to reason about them as shown in the companion technical report.

6 Metatheory

This section establishes the main properties of the system, which are progress and preservation, adequacy, and the most important result: the fundamental theorem where we prove that statically well-typed programs are semantically well-typed. The theorems and their complete proofs are provided in the companion TR [15].

The progress and preservation theorems assume memory locations to be well-formed, $\vdash_{\gamma}^{\mathbf{Md}} \mathbf{NV} \mid \mathbf{V} : \Omega \mid \Sigma$, which is defined similarly to the $\mathbf{NV} \mid \mathbf{V} \Vdash \gamma : \Omega \mid \Sigma$ used in the logical relation, but imposes extra conditions based on the execution mode \mathbf{Md} . It states that γ maps variables in contexts Ω and Σ to the nonvolatile and volatile memories, \mathbf{NV} and \mathbf{V} , respectively, such that their qualifiers and the type of the stored values match. Moreover, it requires specific properties on the contexts depending on \mathbf{Md} ; in atomic mode, each checkpointed location in \mathbf{NV} and Ω must have copies in \mathbf{V} and Σ . We state the theorems below.

Theorem 1 (Progress for Commands). *If $\mathbf{Md} \mid b \mathcal{R} m : \text{nat} \mid \Omega; \Sigma \vdash_{\text{Sig}} c : \tau$, then $\forall n : \text{nat}$ with $n \mathcal{R} m$ and $\forall \gamma, \mathbf{NV}, \mathbf{V}$ with $\vdash_{\gamma}^{\mathbf{Md}} \mathbf{NV} \mid \mathbf{V} : \Omega \mid \Sigma$, either $\gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \mathbf{V} \mid c$ is a value, or for some configuration $\gamma' \mid \mathbf{Md}' \mid n' \mid \mathbf{NV}' \mid \mathbf{V}' \mid c'$ we have $\gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \mathbf{V} \mid c \rightarrow \gamma' \mid \mathbf{Md}' \mid n' \mid \mathbf{NV}' \mid \mathbf{V}' \mid c'$. Moreover, if \mathbf{Md} is an atomic mode, we have $\mathbf{NV}' = \mathbf{NV}$.*

Theorem 2 (Preservation for Commands). *If $\mathbf{Md} \mid b \geq 0 : \text{nat} \mid \Omega; \Sigma \vdash_{\text{Sig}} c : \tau$, and for some $\vdash_{\gamma}^{\mathbf{Md}} \mathbf{NV} \mid \mathbf{V} : \Omega \mid \Sigma$ and $n : \text{nat} \geq 0$, we have $\gamma \mid \mathbf{Md} \mid n \mid \mathbf{NV} \mid \mathbf{V} \mid c \rightarrow \gamma' \mid \mathbf{Md} \mid n' \mid \mathbf{NV}' \mid \mathbf{V}' \mid c'$, then for some Σ_1 , we have $\mathbf{Md} \mid b \geq 0 : \text{nat} \mid \Omega; \Sigma_1 \vdash_{\text{Sig}} c' : \tau$, where $\vdash_{\gamma'}^{\mathbf{Md}} \mathbf{NV}' \mid \mathbf{V}' : \Omega \mid \Sigma_1$ and $n' \geq 0$.*

Theorem 3 (Fundamental Theorem). *If $b : \text{nat} \mid \Omega \vdash p : \uparrow \mathbf{C}_{\text{unit}}$, then $b : \text{nat} \mid \Omega \Vdash p : \uparrow \mathbf{C}_{\text{unit}}$.*

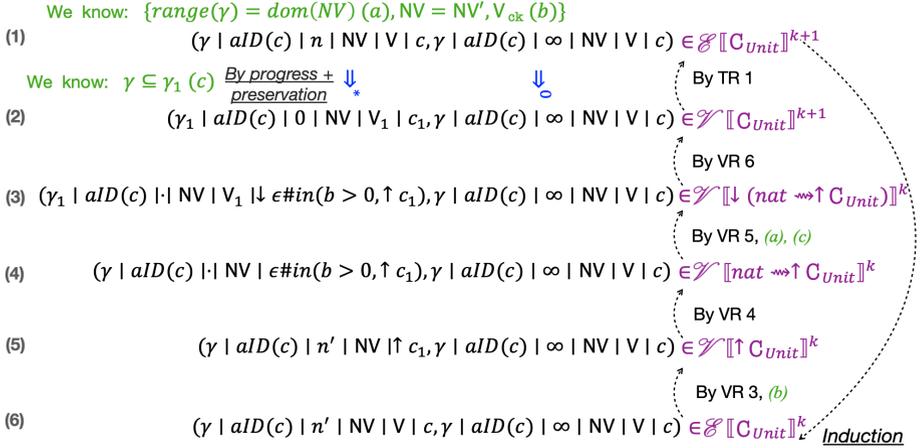


Fig. 13. Proof of the fundamental theorem for P-Ckpt

The proof of Theorem 3 is by induction on the static typing derivation for p and considers the last step in the derivation. Fig. 13 explains the idea of the proof for the case where P-Ckpt is the last step of the derivation. By inversion, $p = \text{Ckpt}[aID, \rho](c); p'$. Also, c is well-typed for static contexts Ω' and Σ , where $\Omega' = \Omega'', \Sigma_{ck}$. The goal is to establish point (1) in the figure: c is related to itself in the term interpretation for arbitrary n, m, γ, NV and V where $NV \mid V \Vdash \gamma :: \Omega'', \Sigma_{ck} \mid \Sigma$. The last condition enforces that the static contexts match the dynamic context. The condition also establishes the more refined well-formedness condition that $\vdash_{\gamma}^{\text{Md}} NV \mid V : \Omega \mid \Sigma$ in atomic mode, required by progress and preservation, since it enforces that each checkpointed location in NV and Ω have copies in V and Σ . In particular, $NV = NV', V_{ck}$ and $range(\gamma) = dom(NV)$. When $m = 0$, the proof is trivial. Consider the case where $m = k + 1$. By the progress and preservation theorems, the first configuration can take multiple steps until it becomes a value $\gamma_1 \mid aID(c) \mid n' \mid NV \mid V_1 \mid c_1$ that continues to be well-typed. If $n' > 0$, the second configuration steps similarly to completion and establishes that the two resulting configurations are in the value relation. This case is not shown in the figure. If $n' = 0$, the second configuration does not step and instead reaches point (2) in Fig. 13. At point (2), the proof must show that the configurations are in the value interpretation at type \mathbb{C}_{Unit} .

The dashed line in the figure states that establishing point (2) implies the relation in point (1). The cascade of implications (dashed lines) follows the definition of the value relations at each type. At each step, we invert on the typing rule of the open configuration and show that runtime memories stay well-defined for static contexts. At point (4), we apply the power failure policy for atomic regions, which drops the volatile memory V_1 and creates a mapping using the domain of NV . By the prior conditions established, we know the created mapping is the original mapping γ . At point (6), we apply the restore policy for atomic regions, which creates a new volatile memory based on NV . Again by the

prior conditions established, we know the volatile memory created is the original volatile V . The goal at point (6) is similar to our original goal at point (1), except that the proof uses an inductive argument to relate the two configurations at k .

Finally the Adequacy Theorem states that semantically well-typed programs are idempotent, defined below. The proof is illustrated in Section 5.2.

Definition 1 (Idempotency). *A triple of a program p , nonvolatile memory NV , and a mapping γ is idempotent, if every intermittent execution of the program can be simulated by a continuous execution of it: for all $n, n', \chi_1, \chi'_1, NV', p'$, if $[\chi_1 \triangleright \varepsilon] \otimes \gamma \mid n \mid NV \mid p \Rightarrow [\chi'_1 \triangleright \varepsilon] \otimes \gamma \mid n' \mid NV' \mid p'$, then $[\chi_2 \triangleright \varepsilon] \otimes \gamma \mid \infty \mid NV \mid p \Rightarrow [\chi_2 \triangleright \varepsilon] \otimes \gamma \mid \infty \mid NV' \mid p'$.*

Theorem 4 (Adequacy). *Consider $b : \text{nat} \mid \Omega \Vdash p : C_{\text{unit}}$, a nonvolatile memory NV and a bijective map γ that matches qualifiers and types from variables in Ω to locations in NV . The triple of p , NV , and γ is idempotent.*

7 Discussion & Related Work

Intermittent Computing. Surbatovich et al. [41] provide the first formal framework for reasoning about intermittent execution, give the correctness definition that we use, and identify precise memory invariants needed for an execution to be correct. Our Crash types capture some of these invariants; capturing all requires reasoning about the effects of non-deterministic sensor inputs, which we leave to future work. This work is the first to treat intermittent operations at the type level and explore the logical interpretation of intermittent execution. We speculate that our type-based approach using logical relations will provide a cleaner foundation for reasoning about the correctness of more complex intermittent systems, e.g., concurrent ones. Other works that investigate the formal properties of intermittent computing either reason about the effects of intermittent execution on peripheral interactions [9] or enforce timeliness constraints on sensor readings [40], which are orthogonal to ours.

Adjoint Logic. Benton et al. [7,8] provided the first categorical foundation for using adjoint functors to combine linear and nonlinear logics and showed that a well-behaved calculus requires an independence principle: linear formulae cannot appear in the assumptions of a nonlinear sequent. Follow up works further generalized the system [20,21,36]. There, the relation to Pfenning and Davies's [30] formulation of the lax \circ modality was noted; \circ corresponds to UF , where F and U are adjunctions between truth and validity categories. Short of a full curry-howard correspondence for our type system and underlying logic, we designed the rules for \uparrow and \downarrow based on the above calculi. Our stable and unstable contexts correspond to the validity and truth contexts respectively. Thus, we speculate that the combination $\uparrow\downarrow$ in our system corresponds to the lax modality.

Several prior works used type systems with adjoint modalities to model switching between program modes [6,14,34], e.g., switching a processes' mode between shared and unshared [6], or adding multicasting, replicable services, and cancellation modes to a session-typed message passing system [34]. We are

the first to use these modalities to handle unforeseen shut-downs and distinguish between stable and power-failure prone modes.

Logical Relations. Prior work [3,42] uses step indexing to ensure the well-foundedness of logical relations that handle heaps with cyclic references, dynamic memory allocation, or recursive types. Our Crash types model the infinite computation that an atomic region can experience under a non-deterministic number of power failures and re-executions. This recursion necessitates an-indexed relation that limits the number of execution attempts a program can make.

Jung and Tiurnyn introduced a logical relation for lambda definability that allows varying arities [18]. The idea is to increase the arity when passing to later worlds instead of starting with a large arity. Our logical relation can also be viewed as a relation with different arities; the initial type of the relation is binary, while after a crash the type of the value relation only corresponds to the intermittent configuration. During these value steps, the relation is unary, with the continuous configuration acting as a kripke world for the intermittent configuration. After restoration, the relation reverts to binary.

Logical relations have been widely used to prove program equivalence, e.g., [2,3,10,16]. At a high level, idempotency is similar to program equivalence, but it handles re-execution and requires us only to prove simulation from an intermittent to continuous run, not vice-versa.

Algebraic Effect Handlers. Algebraic effect handlers [27,31,32,33] give a unified theory for computational effects, e.g., exceptions and interactive input/output. A handler accesses the continuation to transform the computation. Following effect handler syntax, we write effectful environmental interactions of our system as $\varepsilon\#\text{in}(b > 0, \uparrow\kappa)$, where b refers to a natural number returned by the environment and $\uparrow\kappa$ is the continuation. Our restore policy resembles a handler, in that it has access to the continuation, but an atomic region may dismiss the continuation, restarting from a saved command.

Crash Hoare Logic. Crash Hoare logic (CHL) [11] ensures the correctness of crash and restore operations in a file system. CHL extends Hoare logic with a crash condition and a recovery procedure. The crash condition states what happens to the state on a crash. The recovery procedure runs after the crash and manipulates the state before resuming. The system checks that if the program crashes, the storage system will recover to a state consistent with the specifications. Unlike us, they do not care about idempotency, requiring manual effort to formalize the crash condition and recovery policy. Our syntactic typing fixes the power failure, restore, and commit policies, and our formal results guarantee that following the policies ensures idempotency, the common correctness condition for intermittent execution. We also allow the programmer to formalize bespoke semantically well-typed policies.

8 Conclusion

This work provides the first logical interpretation of intermittent execution. It shows that adjoint logic can be applied to define Crash types, which internalize

the dualities between stable and unstable values, and complete versus partial (re-)executions of intermittent programs. The typing constraints capture invariants of power failure, restoration, and re-execution in intermittent systems. The proofs of progress, preservation, and the fundamental theorem imply the correctness of intermittent systems, i.e. idempotency of execution.

References

1. Adkins, J., Campbell, B., Ghena, B., Jackson, N., Pannuto, P., Dutta, P.: The signpost network: Demo abstract. In: Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM. SenSys '16 (2016). <https://doi.org/10.1145/2994551.2996542>
2. Ahmed, A., Dreyer, D., Rossberg, A.: State-dependent representation independence. In: Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. p. 340–353. POPL '09, Association for Computing Machinery, New York, NY, USA (2009). <https://doi.org/10.1145/1480881.1480925>
3. Ahmed, A.J.: Semantics of types for mutable state. Princeton University (2004)
4. Balsamo, D., Weddell, A., Das, A., Arreola, A., Brunelli, D., Al-Hashimi, B., Merrett, G., Benini, L.: Hibernus++: A self-calibrating and adaptive system for transiently-powered embedded devices. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems **PP**(99), 1–1 (2016). <https://doi.org/10.1109/TCAD.2016.2547919>
5. Balsamo, D., Weddell, A.S., Merrett, G.V., Al-Hashimi, B.M., Brunelli, D., Benini, L.: Hibernus: Sustaining computation during intermittent supply for energy-harvesting systems. IEEE Embedded Systems Letters **7**(1), 15–18 (2015). <https://doi.org/10.1109/LES.2014.2371494>
6. Balzer, S., Toninho, B., Pfenning, F.: Manifest deadlock-freedom for shared session types. In: Proceedings of the 29th European Symposium on Programming. pp. 611–639 (2019)
7. Benton, N., Wadler, P.: Linear logic, monads and the lambda calculus. In: Proceedings 11th Annual IEEE Symposium on Logic in Computer Science. pp. 420–431. IEEE (1996)
8. Benton, P.N.: A mixed linear and non-linear logic: Proofs, terms and models. In: International Workshop on Computer Science Logic. pp. 121–135. Springer (1994)
9. Berthou, G., Dagand, P.E., Demange, D., Oudin, R., Risset, T.: Intermittent computing with peripherals, formally verified. In: The 21st ACM SIGPLAN/SIGBED Conference on Languages, Compilers, and Tools for Embedded Systems. pp. 85–96. LCTES '20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3372799.3394365>
10. Birkedal, L., Støvring, K., Thamsborg, J.: Realizability semantics of parametric polymorphism, general references, and recursive types. In: International Conference on Foundations of Software Science and Computational Structures. pp. 456–470. FOSSACS '09, Springer (2009)
11. Chen, H., Ziegler, D., Chajed, T., Chlipala, A., Kaashoek, M.F., Zeldovich, N.: Using crash hoare logic for certifying the fscq file system. In: Proceedings of the 25th Symposium on Operating Systems Principles. pp. 18–37. SOSP '15 (2015)

12. Colin, A., Lucia, B.: Chain: Tasks and channels for reliable intermittent programs. In: Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications. OOPSLA '16 (2016). <https://doi.org/10.1145/2983990.2983995>
13. Dahiya, M., Bansal, S.: Automatic verification of intermittent systems. In: Dillig, I., Palsberg, J. (eds.) Verification, Model Checking, and Abstract Interpretation. VMCAI '18 (2018)
14. Das, A., Balzer, S., Hoffmann, J., Pfenning, F., Santurkar, I.: Resource-aware session types for digital contracts. In: IEEE 34th Computer Security Foundations Symposium. pp. 1–16. CSF '21 (2021)
15. Derakhshan, F., Dotzel, M., Surbatovich, M., Jia, L.: Technical report: Modal crash types for intermittent computing. Tech. rep., Carnegie Mellon University (2023). <https://doi.org/10.1184/R1/21950804>
16. Dreyer, D., Neis, G., Birkedal, L.: The impact of higher-order state and control effects on local relational reasoning. *Journal of Functional Programming* **22**(4-5), 477–528 (2012)
17. Hester, J., Storer, K., Sorber, J.: Timely execution on intermittently powered batteryless sensors. In: Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems (2017). <https://doi.org/10.1145/3131672.3131673>
18. Jung, A., Tiuryn, J.: A new characterization of lambda definability. In: International Conference on Typed Lambda Calculi and Applications. pp. 245–257. Springer (1993)
19. Kortbeek, V., Yildirim, K.S., Bakar, A., Sorber, J., Hester, J., Pawelczak, P.: Time-sensitive intermittent computing meets legacy software. In: Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems. pp. 85–99. ASPLOS '20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3373376.3378476>
20. Licata, D.R., Shulman, M.: Adjoint logic with a 2-category of modes. In: International Symposium on Logical Foundations of Computer Science. pp. 219–235. Springer (2016)
21. Licata, D.R., Shulman, M., Riley, M.: A fibrational framework for substructural and modal logics. In: 2nd International Conference on Formal Structures for Computation and Deduction. FSCD '17, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2017)
22. Lucia, B., Denby, B., Manchester, Z., Desai, H., Ruppel, E., Colin, A.: Computational nanosatellite constellations: Opportunities and challenges. *GetMobile: Mobile Comp. and Comm.* **25**(1), 16–23 (Jun 2021). <https://doi.org/10.1145/3471440.3471446>
23. Lucia, B., Ransford, B.: A simpler, safer programming and execution model for intermittent systems. In: Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation. PLDI '15 (2015). <https://doi.org/10.1145/2737924.2737978>
24. Maeng, K., Colin, A., Lucia, B.: Alpaca: Intermittent execution without checkpoints. *Proc. ACM Program. Lang.* **1**(OOPSLA), 96:1–96:30 (Oct 2017). <https://doi.org/10.1145/3133920>
25. Maeng, K., Lucia, B.: Supporting peripherals in intermittent systems with just-in-time checkpoints. In: Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation. p. 1101–1116. PLDI '19 (2019). <https://doi.org/10.1145/3314221.3314613>

26. Maeng, K., Lucia, B.: Adaptive low-overhead scheduling for periodic and reactive intermittent execution. In: Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation. pp. 1005–1021. PLDI '20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3385412.3385998>
27. Moggi, E.: Computational lambda-calculus and monads. University of Edinburgh, Department of Computer Science, Laboratory for Foundations of Computer Science (1988)
28. Nardello, M., Desai, H., Brunelli, D., Lucia, B.: Camaroptera: A batteryless long-range remote visual sensing system. In: Proceedings of the 7th International Workshop on Energy Harvesting & Energy-Neutral Sensing Systems. pp. 8–14. ENSys'19, ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3362053.3363491>
29. NASA: What is KickSat-2? <https://www.nasa.gov/ames/kicksat> (2019), visited April 15th, 2022
30. Pfenning, F., Davies, R.: A judgmental reconstruction of modal logic. *Mathematical structures in computer science* **11**(4), 511–540 (2001)
31. Plotkin, G., Power, J.: Semantics for algebraic operations. *Electronic Notes in Theoretical Computer Science* **45**, 332–345 (2001)
32. Plotkin, G., Pretnar, M.: Handlers of algebraic effects. In: Proceedings of the 19th European Symposium on Programming. pp. 80–94. Springer (2009)
33. Pretnar, M., Plotkin, G.D.: Handling algebraic effects. *Logical methods in computer science* **9** (2013)
34. Pruiksma, K., Pfenning, F.: A message-passing interpretation of adjoint logic. *Journal of Logical and Algebraic Methods in Programming* **120**, 100637 (2021)
35. Ransford, B., Sorber, J., Fu, K.: Mementos: System support for long-running computation on RFID-scale devices. In: Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems. ASPLOS XVI (2011). <https://doi.org/10.1145/1950365.1950386>
36. Reed, J.: A judgmental deconstruction of modal logic. Unpublished manuscript, January (2009)
37. Ruppel, E., Lucia, B.: Transactional concurrency control for intermittent, energy-harvesting computing systems. In: Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation. p. 1085–1100. PLDI '19 (2019). <https://doi.org/10.1145/3314221.3314583>
38. Shavit, N., Touitou, D.: Software transactional memory. In: Proceedings of the fourteenth annual ACM symposium on Principles of distributed computing. pp. 204–213. PODC '95 (1995). <https://doi.org/10.1145/224964.224987>
39. Surbatovich, M., Jia, L., Lucia, B.: I/o dependent idempotence bugs in intermittent systems. *Proc. ACM Program. Lang.* **3**(OOPSLA), 183:1–183:31 (Oct 2019). <https://doi.org/10.1145/3360609>
40. Surbatovich, M., Jia, L., Lucia, B.: Automatically enforcing fresh and consistent inputs in intermittent systems. In: Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation. p. 851–866. PLDI '21, Association for Computing Machinery, New York, NY, USA (2021). <https://doi.org/10.1145/3453483.3454081>
41. Surbatovich, M., Lucia, B., Jia, L.: Towards a formal foundation of intermittent computing. *Proc. ACM Program. Lang.* **4**(OOPSLA) (Nov 2020). <https://doi.org/10.1145/3428231>
42. Thamsborg, J., Birkedal, L.: A kripke logical relation for effect-based program transformations. *ACM SIGPLAN Notices* **46**(9), 445–456 (2011)

43. Van Der Woude, J., Hicks, M.: Intermittent computation without hardware support or programmer intervention. In: Proceedings of OSDI'16: 12th USENIX Symposium on Operating Systems Design and Implementation (2016). <https://doi.org/10.5555/3026877.3026880>
44. Yildirim, K.S., Majid, A.Y., Patoukas, D., Schaper, K., Pawelczak, P., Hester, J.: Ink: Reactive kernel for tiny batteryless sensors. In: Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems. pp. 41–53. SenSys '18, ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3274783.3274837>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

