

Pipit on the Post: Proving Pre- and Post-conditions of Reactive Systems

Amos Robinson ✉ 

Australian National University, Canberra, Australia

Alex Potanin ✉ 

Australian National University, Canberra, Australia

Abstract

Reactive languages such as Lustre and Scade are used to implement safety-critical control systems; proving such programs correct and having the proved properties apply to the compiled code is therefore equally critical. We introduce Pipit, a small reactive language embedded in F^* , designed for verifying control systems and executing them in real-time. Pipit includes a verified translation to transition systems; by reusing F^* 's existing proof automation, certain safety properties can be automatically proved by k-induction on the transition system. Pipit can also generate executable code in a subset of F^* which is suitable for compilation and real-time execution on embedded devices. The executable code is deterministic and total and preserves the semantics of the original program.

2012 ACM Subject Classification Computer systems organization → Real-time languages; Theory of computation → Program verification; Software and its engineering → Specialized application languages

Keywords and phrases Lustre, streaming, reactive, verification

1 Introduction

Safety-critical control systems, such as the anti-lock braking systems that are present in most cars today, need to be correct and execute in real-time. One approach, favoured by parts of the aerospace industry, is to implement the controllers in a high-level language such as Lustre [10] or Scade [13], and verify that the implementations satisfy the high-level specification using a model-checker, such as Kind2 [11]. These model-checkers can prove many interesting safety properties automatically, but do not provide many options for manual proofs when the automated proof techniques fail. Additionally, the semantics used by the model-checker may not match the semantics of the compiled code, in which case properties proved do not necessarily hold on the real system. This mismatch may occur even when the compiler has been verified to be correct, as in the case of Vélus [5]. For example, in Vélus, integer division rounds towards zero, matching the semantics of C; however, integer division in Kind2 rounds to negative infinity, matching SMT-lib [2, 25].

To be confident that our proofs hold on the real system, we need a single semantics that is shared between the compiler and the model-checker or prover. In this paper we introduce Pipit¹, an embedded domain-specific language for implementing and verifying controllers in F^* . Pipit aims to provide a high-level language based on Lustre, while reusing F^* 's proof automation and manual proofs for verifying controllers [31], and using Low*'s C-code generation for real-time execution [34]. To verify programs, Pipit translates its expression language to a transition system for k-inductive proofs, which is verified to be an abstraction of the original semantics. To execute programs, Pipit can generate executable code, which is total and semantics-preserving.

In this paper, we make the following contributions:

¹ Implementation available at <https://github.com/songlarknet/pipit>

- 43 ■ we motivate the need to combine manual and automated proofs of reactive systems with
- 44 a strong specification language (Section 2);
- 45 ■ we introduce Pipit, a minimal reactive language that supports rely-guarantee contracts
- 46 and properties; crucially, proof obligations are annotated with a status — *valid* or *deferred*
- 47 — allowing proofs to be delayed until more is known of the program context (Section 3);
- 48 ■ we describe a *checked semantics* for Pipit, which is parameterised by the property status;
- 49 after checking deferred properties, programs can be *blessed*, and their properties lifted to
- 50 valid status (Subsection 3.2);
- 51 ■ we describe an encoding of transition systems that can express under-specified rely-
- 52 guarantee contracts as functions rather than relations; composing functions results in
- 53 simpler transition systems (Section 4);
- 54 ■ we identify the invariants and lemmas required to prove that the abstract transition
- 55 system is an abstraction of the original semantics (Subsection 3.3, Subsection 4.1);
- 56 ■ similarly, we offer a mechanised proof that the executable transition system preserves the
- 57 original semantics (Section 5);
- 58 ■ finally, we evaluate Pipit by implementing the high-level logic of a time-triggered Controller
- 59 Area Network (CAN) bus driver, which we have partially verified (Section 6).

60 2 Pipit for time-triggered networks

61 To introduce Pipit, we consider a driver with a static schedule of *triggers*, or actions to be
 62 performed at a particular time; this driver is a simplification of the time-triggered Controller
 63 Area Network (CAN) bus specification [15] which we will discuss further in Section 6.

64 2.1 Deferring and proving properties

65 The schedule of our time-triggered driver is determined by a constant array of triggers, sorted
 66 by their associated time-mark. The driver maintains an index that refers to the current
 67 trigger. At each instant in time, the driver checks if the current trigger has expired or
 68 is inactive, and if so, it increments the index. We first implement a streaming function
 69 *count_when* to maintain the index; the function takes a constant natural number *max* and a
 70 stream of booleans *inc*. At each time step, *count_when* checks whether the current increment
 71 flag is true; if so, it increments the previous counter, saturating at the maximum; otherwise,
 72 it leaves the previous counter as-is.

```

let count_when (max: ℕ) (inc: stream ℬ): stream ℕ =
  rec count
    check? (0 ≤ count ≤ max);
    let count' = (0 fby count) + (if inc then 1 else 0) in
    if count' ≥ max then max else count'

```

73 The implementation of *count_when* first defines a recursive stream, *count*, which states
 74 an invariant about the count before defining the incremented stream *count'*. Inside *count'*,
 75 the syntax *0 fby count* is read as “the initial value of zero *followed by* the previous count”.

76 The syntax *check_? (0 ≤ count ≤ max)* asserts that the count is within the range $[0, max]$.
 77 The subscript *?* on the check is the *property status*, which in this case denotes that the
 78 assertion has been stated, but it is not yet known whether it holds. A property status of
 79 *✓*, on the other hand, denotes that a property has been proved to hold. These property
 80 statuses are used to defer checking properties until enough is known about the environment,

81 and to avoid rechecking properties that have already been proven. In practice, the user
 82 does not explicitly specify property statuses in the source language. The stated property
 83 ($0 \leq count \leq max$) is a stream of booleans which must always be true. Non-streaming
 84 operations such as \leq are implicitly lifted to streaming operations, and non-streaming values
 85 such as 0 and max are implicitly lifted to constant streams.

86 We defer the proof of the property here because, at the point of stating the property
 87 inside the `rec` combinator, we don't yet have a concrete definition for the count variable.
 88 In this case, we could have instead deferred the *statement* of the property by introducing
 89 a let-binding for the recursive count and putting the `check` outside of the `rec` combinator.
 90 However, it is not always possible to defer property statements: for example, when calling
 91 other streaming functions that have their own preconditions, it may not be possible to move
 92 the function call outside of its enclosing `rec`.

93 Pipit is an embedded domain-specific language. The program above is really syntactic
 94 sugar for an F^* program that takes a natural number and constructs a Pipit core expression
 95 with a free boolean variable. We will discuss the details of the core language in Section 3,
 96 but for now we focus on the source program with some minor embedding details omitted.

97 To actually prove the property above, we use the meta-language F^* 's tactics to translate
 98 the program into a transition system and prove the property inductively on the system.
 99 Finally, we *bless* the expression, which marks the properties as valid ($[?] := \checkmark$). Blessing is
 100 an intensional operation: it traverses the expression and updates the internal metadata, but
 101 it does not affect the runtime semantics.

```

let count_when $\checkmark$  (max: N): stream B → stream N =
  let system = System.translate1(count_when max) in
  assert (System.inductive_check system) by (pipit_simplify ());
  bless1 (count_when max)

```

102 The subscript 1 in the translation to transition system and blessing operations refers
 103 to the fact that the stream function has one stream parameter. The *pipit_simplify* tactic
 104 in the assertion performs normalisation-by-evaluation to simplify away the translation to a
 105 first-order transition system; F^* 's proof-by-SMT can then solve the inductive check directly.

106 Callers of *count_when* can now use the validated variant without needing to re-prove
 107 the count-range property. In a dedicated model-checker such as Kind2 [11] or Lesar [35],
 108 this kind of bookkeeping would all be performed under-the-hood. By embedding Pipit in a
 109 general-purpose theorem prover, we move some of the bookkeeping burden onto the user;
 110 however, we have increased confidence that the compiled code matches the verified code and,
 111 as we shall see, we also have access to a rich specification language.

112 2.2 The time-triggered system matrix

113 The schedule of the time-triggered network is abstractly described by a *system matrix*,
 114 consisting of rows of *basic cycles*, columns of *transmission columns*, and cells of optional
 115 messages. Each basic cycle is identified by its cycle index and each transmission column has
 116 an associated time-mark.

117 Figure 1 (left) shows an example system matrix with cycles C0 and C1 and transmission
 118 columns at time-marks 0, 1 and 2. For this example, we assume that one message can be sent
 119 per clock cycle. To execute this system matrix, we synchronise the local time to zero at the
 120 start of basic cycle C0. After a basic cycle completes, the nodes on the network synchronise
 121 before execution continues to the next basic cycle.

| | TM0 | TM1 | TM2 | |
|----|-------|-------|-------|---|
| C0 | MSG A | MSG B | - | 0: { time_mark = 0; enabled = {C0,C1}; msg = A; } |
| C1 | MSG A | - | MSG C | 1: { time_mark = 1; enabled = {C0}; msg = B; } |
| | | | | 2: { time_mark = 2; enabled = {C1}; msg = C; } |

■ **Figure 1** Left: system matrix; right: corresponding triggers array configuration

122 Figure 1 (right) shows the corresponding configuration for the triggers array. The enabled
123 set denotes the basic cycles for which a trigger is active.

124 The system has strict timing requirements which restrict how triggers can be defined. In
125 this example, each trigger has a unique time; in general, trigger times can overlap, but they
126 need to be enabled on distinct cycles. Additionally, the schedule must allow sufficient time
127 for the driver to skip over the disabled triggers. Concretely, we could postpone trigger 1 to
128 send message B at time-mark 2, as triggers 1 and 2 have distinct cycles. However, we could
129 not bring forward trigger 2 to send message C at time-mark 1: the driver can only process
130 one trigger per tick, and it takes two steps to reach trigger 2 from the start of the array.

131 We impose three main restrictions on the triggers array: the time-marks must be sorted;
132 there must be an adequate time-gap between any two triggers that are enabled on the same
133 cycle index; and each trigger's time-mark must be greater-than-or-equal to its index.

134 With these restrictions in place, we prove a lemma *lemma_can_reach_next*, which states
135 that for all valid cycle indices and trigger indices, if the current trigger is enabled in the
136 current cycle and there is another enabled trigger scheduled to occur somewhere in the array
137 after the current one, then there is an adequate time-gap to allow the driver to skip over any
138 disabled triggers in-between. These properties are straightforward in a theorem prover, but
139 would be difficult to state in a model-checker with a limited specification language.

140 2.3 Instantiating lemmas and defining contracts

141 We can now implement the trigger-fetch logic, which keeps track of the current trigger. The
142 trigger-fetch logic uses the *count_when* streaming function to define the index of the current
143 trigger; we tell *count_when* to increment the index whenever the previous index has expired
144 or is inactive in the current basic cycle. We simplify our presentation here and only consider
145 a single cycle in isolation: the real system presented in Section 6 has some extra complexity
146 such as resetting the index, incrementing the cycle index at the start of a new cycle, and
147 using machine integers.

```

let trigger_fetch (cycle: ℕ) (time: stream ℕ): stream ℕ =
  rec index.
    let inc = false fby ((time_mark index) ≤ time ∨ ¬(enabled index cycle)) in
    let index = count_when□ trigger_count inc in
    pose1 (lemma_can_reach_next cycle) index;
    check□ (can_reach_next_active cycle time index);
    index

```

148 The *trigger_fetch* function takes a static cycle index and a stream denoting the current
149 time. The increment flag and the index are mutually dependent — the increment flag depends
150 on the previous value of the index, while the index depends on the current value of the
151 increment flag — so we introduce a recursive stream for the index. We allow the index to go
152 one past the end of the array to denote that there are no more triggers.

153 We use the *pose₁* helper function to lift the *lemma_can_reach_next* lemma to a streaming
154 context and instantiate it; the subscript 1 indicates that the lemma is being applied to

155 one streaming argument (the index). We then state an invariant as a deferred property.
 156 Informally, the invariant states that, either the current active trigger is not late, or the next
 157 active trigger after the current index is in the future and we can reach it in time.

158 With the explicitly instantiated lemma, we can prove the streaming invariant by straight-
 159 forward induction on the transition system. To help compose this function with the rest of
 160 the system, we also abstract over the details of the trigger-fetch mechanism by introducing a
 161 rely-guarantee contract for *trigger_fetch*. The contract we state is that if the environment
 162 ensures that the time doesn't skip — that is, we are called once per microsecond — then we
 163 guarantee that we never encounter a late trigger.

```

let trigger_fetch $\square$  (cycle:  $\mathbb{N}$ ): stream  $\mathbb{N}$   $\rightarrow$  stream  $\mathbb{N}$  =
  let contract = Contract.contract_of_stream1 {
    rely = ( $\lambda$ time. time_no_skips time)
    guar = ( $\lambda$ time index. (index_valid index  $\wedge$  enabled index cycle)
            $\implies$  (time_mark index)  $\geq$  time)
    body = ( $\lambda$ time. trigger_fetch cycle time)
  } in
  assert (Contract.inductive_check contract) by (pipit_simplify ());
  Contract.stream_of_contract1 contract

```

164 In the implementation of the validated variant of *trigger_fetch*, we first construct the
 165 contract from streaming functions. The `Contract.contract_of_stream1` combinator describes
 166 a contract with one input (the time stream), and takes stream transformers for each of the
 167 rely, guarantee and body. The combinator transforms the surface syntax into core expressions.
 168 The assertion (`Contract.inductive_check contract`) then translates the expressions into a
 169 transition system, and checks that if the rely always holds then the guarantee always holds,
 170 and that the as-yet-unchecked subproperties hold. Finally, `Contract.stream_of_contract1`
 171 blesses the core expression and converts it back to a stream transformer, so it can be easily
 172 used by other parts of the program.

173 When this function is used in other parts of the program, the caller must ensure that
 174 the environment satisfies the rely clause. In the core language, this is tracked by another
 175 deferred property status attached to the contract; we will discuss this further in Section 3.

176 **3 Core language**

177 We now introduce the core Pipit language. Note that this form differs slightly from the
 178 surface syntax presented earlier in Section 2, which used the syntax of the metalanguage F^* ,
 179 as well as including proofs in F^* itself.

180 Figure 2 defines the grammar of Pipit. The expression form e includes standard syntax for
 181 values (v), variables (x) and primitive applications ($p(\bar{e})$). Most of the expression forms were
 182 introduced informally in Section 2 and correspond to the clock-free expressions of Lustre [10].

183 The expression syntax for delayed streams (v **fb** e) denotes the previous value of the
 184 stream e , with an initial value of v when there is no previous value.

185 Recursive streams, which can refer to previous values of the stream itself, are defined using
 186 the fixpoint operator (**rec** x . $e[x]$); the syntax $e[x]$ means that the variable x can occur in e .
 187 As in Lustre, recursive streams can only refer to their previous values and must be *guarded*
 188 by a delay: the stream (**rec** x . 0 **fb** $(x + 1)$) is well-defined, but stream (**rec** x . $x + 1$) is
 189 invalid and has no computational interpretation. This form of recursion differs slightly from
 190 standard Lustre, which uses a set of mutually-recursive bindings. Although we cannot express

| | |
|---|---|
| $ \begin{aligned} e, e' &:= v \mid x \mid p(\bar{e}) \\ &\mid v \text{ fby } e \mid \text{rec } x. e[x] \\ &\mid \text{let } x = e \text{ in } e'[x] \\ &\mid \text{check}_\pi e_{\text{prop}} \\ &\mid \text{contract}_\pi \{e_{\text{rely}}\} e_{\text{body}} \{x. e_{\text{guar}}[x]\} \end{aligned} $ | <p>(values, variables and operations)</p> <p>(delayed and recursive streams)</p> <p>(let-expressions)</p> <p>(checked properties)</p> <p>(rely-guarantee contracts)</p> |
| $ \begin{aligned} v &:= n \in \mathbb{N} \mid b \in \mathbb{B} \mid r \in \mathbb{R} \mid \dots \\ p &:= (+) \mid (-) \mid (\times) \mid \text{if-then-else} \mid \dots \end{aligned} $ | <p>(values)</p> <p>(primitives)</p> |
| $\pi := \square \mid \boxminus$ | <p>(property statuses: valid or unknown)</p> |
| $ \begin{aligned} V &:= \cdot \mid V; v \\ \sigma &:= \{\bar{x} \mapsto \bar{v}\} \\ \Sigma &:= \cdot \mid \Sigma; \sigma \\ \tau, \tau' &:= \mathbb{N} \mid \mathbb{B} \mid \tau \times \tau \mid \dots \\ \Gamma &:= \cdot \mid x : \tau, \Gamma \end{aligned} $ | <p>(streams of values)</p> <p>(heaps)</p> <p>(streaming history environments)</p> <p>(value types)</p> <p>(type environments)</p> |

■ **Figure 2** Pipit core language grammar, which contains expressions e , values v , primitive operations p , and property statuses π .

191 mutually-recursive bindings in the core syntax here, we can express them as a notation on
 192 the surface syntax by combining the bindings together into a single tuple.

193 Checked properties and contracts are annotated with their property status π , which can
 194 either be valid (\square) or unknown (\boxminus). For checked properties $\text{check}_\pi e$, the property status
 195 denotes whether the property has been proved to be valid.

196 Contracts $\text{contract}_\pi \{e_{\text{rely}}\} e_{\text{body}} \{x. e_{\text{guar}}[x]\}$ involve two verification conditions. Firstly,
 197 when a contract is *defined*, the definer must prove that the body e_{body} satisfies the contract:
 198 roughly, if e_{rely} is always true, then $e_{\text{guar}}[x := e_{\text{body}}]$ is always true. Secondly, when a contract
 199 is *instantiated*, the caller must prove that the environment satisfies the precondition: that is,
 200 e_{rely} is always true. Conceptually, then, a contract could have two property statuses: one for
 201 the definer, and one for the instantiation. However, in practice, it is not useful to defer the
 202 proof of a contract definition — one could achieve a similar effect by replacing the contract
 203 with its implementation. For this reason, we only annotate contracts with one property
 204 status, which denotes whether the instantiation has been proved to satisfy the precondition.

205 Streams V are represented as a sequence of values; streaming history environments Σ are
 206 streams of heaps. Types τ and type environments Γ are standard.

207 We define the typing judgments for Pipit in Figure 3. Most of the typing rules are standard
 208 for an unlocked Lustre. The typing judgment $\Gamma \vdash e : \tau$ denotes that, in an environment
 209 of streams Γ , expression e denotes a stream of type τ . This core typing judgment differs
 210 from the surface syntax used in Section 2, which used an explicit stream type; for the core
 211 language, we instead assume that everything is a stream.

212 For values, we use an auxiliary judgment form $\text{prim-value-type}(v) = \tau$ to denote that value
 213 v has type τ . Likewise, for primitives we use the auxiliary judgment form $\text{prim-type}(p) =$
 214 $(\tau_1 \times \dots \times \tau_n) \rightarrow \tau'$ to denote that primitive p takes arguments of type τ_i and returns a
 215 result of type τ' . Primitives are pure, non-streaming functions.

216 Rules TVALUE, TVAR, TPRIM and TLET are standard.

217 Rule TFBY states that expression $v \text{ fby } e$ requires both v and e to have equal types; the

$$\boxed{\Gamma \vdash e : \tau}$$

$$\begin{array}{c} \frac{\text{prim-value-type}(v) = \tau}{\Gamma \vdash v : \tau} \text{ (TVALUE)} \qquad \frac{}{\Gamma, x : \tau, \Gamma' \vdash x : \tau} \text{ (TVAR)} \\ \\ \frac{\text{prim-type}(p) = (\tau_1 \times \dots \times \tau_n) \rightarrow \tau' \quad \Gamma \vdash e_1 : \tau_1 \quad \dots \quad \Gamma \vdash e_n : \tau_n}{\Gamma \vdash p(\bar{e}) : \tau'} \text{ (TPRIM)} \\ \\ \frac{\text{prim-value-type}(v) = \tau \quad \Gamma \vdash e' : \tau}{\Gamma \vdash v \text{ fby } e' : \tau} \text{ (TFBY)} \qquad \frac{\Gamma, x : \tau \vdash e : \tau}{\Gamma \vdash \text{rec } x. e[x] : \tau} \text{ (TREC)} \\ \\ \frac{\Gamma \vdash e : \tau \quad \Gamma, x : \tau \vdash e' : \tau'}{\Gamma \vdash \text{let } x = e \text{ in } e'[x] : \tau'} \text{ (TLET)} \qquad \frac{\Gamma \vdash e : \mathbb{B}}{\Gamma \vdash \text{check}_\pi e : \text{unit}} \text{ (TCHECK)} \\ \\ \frac{\Gamma \vdash e_{\text{rely}} : \mathbb{B} \quad \Gamma \vdash e_{\text{body}} : \tau \quad \Gamma, x : \tau \vdash e_{\text{guar}} : \mathbb{B}}{\Gamma \vdash \text{contract}_\pi \{e_{\text{rely}}\} e_{\text{body}} \{x. e_{\text{guar}}[x]\} : \tau} \text{ (TCONTRACT)} \end{array}$$

■ **Figure 3** Typing rules for Pipit; the judgment $\Gamma \vdash e : \tau$ denotes that expression e describes a *stream* of values of type τ . Two auxiliary judgment forms are used for values and primitive operations; their rules are standard and are omitted.

218 result is the same type.

219 Rule TREC states that a recursive stream $\text{rec } x. e$ has the recursive stream bound inside
220 e . The recursion must also be guarded, in that any recursive references to x are delayed, but
221 this requirement is performed as a separate syntactic check described in Subsection 3.3.

222 Rule TCHECK states that statically checking a property $\text{check}_\pi e$ requires a boolean
223 property e and returns unit.

224 Finally, rule TCONTRACT applies for a contract $\text{contract}_\pi \{e_{\text{rely}}\} e_{\text{body}} \{x. e_{\text{guar}}[x]\}$
225 with a body expression of some type τ . The overall expression has result type τ . Both rely
226 and guarantee clauses must be boolean expressions. Additionally, the guarantee clause can
227 refer to the result value by x .

228 3.1 Dynamic semantics

229 The dynamic semantics of Pipit are defined in Figure 4. We present our semantics in a
230 big-step form. This differs somewhat from traditional *reactive* semantics of Lustre [10]. Our
231 big-step semantics emphasises the equational nature of Pipit, as it is substitution-based,
232 while the reactive semantics emphasises the finite-state streaming execution of the system.
233 We use transition systems for reasoning about the finite-state execution (Section 4), which is
234 fairly standard [9, 11, 35]. Previous work on the W-CALCULUS [17] for linear digital signal
235 processing filters makes a similar distinction and provides a non-streaming semantics for
236 reasoning about programs and a streaming semantics for executing programs.

237 The judgment form $\Sigma \vdash e \Downarrow v$ denotes that expression e evaluates to value v under
238 streaming history Σ . The streaming history is a stream of heaps; in practice, we only evaluate
239 expressions with a non-empty streaming history.

240 Rule VALUE states that evaluating a value results in the value itself.

$$\boxed{\Sigma \vdash e \Downarrow v}$$

$$\frac{}{\Sigma \vdash v \Downarrow v} \text{ (VALUE)} \quad \frac{}{\Sigma; \sigma \vdash x \Downarrow \sigma(x)} \text{ (VAR)}$$

$$\frac{\Sigma \vdash e_1 \Downarrow v_1 \quad \dots \quad \Sigma \vdash e_n \Downarrow v_n}{\Sigma \vdash p(\bar{e}) \Downarrow \text{prim-sem}(p, \bar{v})} \text{ (PRIM)}$$

$$\frac{}{\sigma \vdash v \text{ fby } e' \Downarrow v} \text{ (FBY}_1\text{)} \quad \frac{\text{length}(\Sigma) > 0 \quad \Sigma \vdash e' \Downarrow v'}{\Sigma; \sigma \vdash v \text{ fby } e' \Downarrow v'} \text{ (FBY}_S\text{)}$$

$$\frac{\Sigma \vdash e[x := \text{rec } x. e] \Downarrow v}{\Sigma \vdash \text{rec } x. e[x] \Downarrow v} \text{ (REC)} \quad \frac{\Sigma \vdash e'[x := e] \Downarrow v}{\Sigma \vdash \text{let } x = e \text{ in } e'[x] \Downarrow v} \text{ (LET)}$$

$$\frac{}{\Sigma \vdash \text{check}_\pi e \Downarrow ()} \text{ (CHECK)}$$

$$\frac{\Sigma \vdash e_{\text{body}} \Downarrow v}{\Sigma \vdash \text{contract}_\pi \{e_{\text{rely}}\} e_{\text{body}} \{x. e_{\text{guar}}[x]\} \Downarrow v} \text{ (CONTRACT)}$$

$$\boxed{\Sigma \vdash e \Downarrow^* V} \quad \boxed{\Sigma \vdash e \Downarrow^\square \top}$$

$$\frac{}{\cdot \vdash e \Downarrow^* \cdot} \text{ (STEPS}_0\text{)} \quad \frac{\Sigma \vdash e \Downarrow V \quad \Sigma; \sigma \vdash e \Downarrow v}{\Sigma; \sigma \vdash e \Downarrow V; v} \text{ (STEPS}_S\text{)}$$

$$\frac{\Sigma \vdash e \Downarrow^* \top; \dots}{\Sigma \vdash e \Downarrow^\square \top} \text{ (ALWAYS)}$$

■ **Figure 4** Dynamic semantics for Pipit; the judgment form $\Sigma \vdash e \Downarrow v$ denotes that evaluating expression e under streaming history Σ results in value v .

241 Rule VAR states that to evaluate a variable x under some non-empty stream history $\Sigma; \sigma$,
 242 where σ is the most recent heap, we look up the variable in σ .

243 Rule PRIM states that to evaluate a primitive p applied to many arguments e_1 to e_n , we
 244 evaluate each argument separately; we then use the prim-sem metafunction to apply the
 245 primitive.

246 Rule FBY₁ evaluates a followed-by expression when the streaming history contains only a
 247 single element. Here, $v \text{ fby } e$ evaluates to v , as there is no previous value of e to use.

248 Rule FBY_S evaluates a followed-by expression when the streaming history contains
 249 multiple entries. In this case, $v \text{ fby } e$ proceeds to evaluate the previous value of e by
 250 discarding the most recent entry from the streaming history.

251 Rule REC evaluates a recursive stream $\text{rec } x. e$ by unfolding the recursion one step. For
 252 causal expressions (Subsection 3.3), where each recursive occurrence of x is guarded by a
 253 followed-by, this unfolding will eventually terminate, as the follow-by shortens the streaming
 254 history.

255 Rule LET is standard.

256 Rule CHECK states that check expressions always evaluate to unit. We do not perform a
 257 dynamic check that the property is true here; checking the truth of properties is dealt with
 258 in the checked semantics (Subsection 3.2).

259 Rule CONTRACT states that contracts evaluate by just evaluating their body. Like with
 260 checks, we do not perform a dynamic check that the precondition and postcondition hold.

261 We also use two auxiliary judgment forms: $\Sigma \vdash e \Downarrow^* V$ and $\Sigma \vdash e \Downarrow^\square \top$.

262 Judgment form $\Sigma \vdash e \Downarrow^* V$ denotes that, under streaming history Σ , expression e
 263 evaluates to the *stream* V . This judgment is an iterated application of the single-value
 264 big-step form.

265 Judgment form $\Sigma \vdash e \Downarrow^\square \top$ denotes that expression e , which must be a boolean, evaluates
 266 to the stream of trues under history Σ . Informally, it can be read as “in streaming history Σ ,
 267 e is always true”.

268 3.2 Checked semantics

269 In addition to the big-step semantics above, we also define a judgment form for checking that
 270 the properties and contracts of a program hold for a particular streaming history. We call
 271 these the *checked* semantics. Unlike an axiomatic semantics, the checked semantics operate
 272 on a concrete set of input streams.

273 The checked semantics have the judgment form $\Sigma \vdash_\pi e$ valid, which denotes that under
 274 streaming history Σ , the properties of e with status π hold. The property status dictates
 275 which properties should be checked and which should be ignored.

276 To show that an expression e 's unknown properties hold, we prove that for all streaming
 277 histories Σ , assuming the valid properties hold ($\Sigma \vdash_{\square} e$ valid), then the unknown properties
 278 ($\Sigma \vdash_{\square} e$ valid) hold. The assumption here means that we do not have to re-check properties
 279 after proving them once.

280 Contracts involve two proofs: one for the definition and one for the instantiation. To prove
 281 that a contract definition $\text{contract}_\pi \{e_{\text{rely}}\} e_{\text{body}} \{x. e_{\text{guar}}[x]\}$ is valid, we show that for all
 282 streaming histories Σ , assuming the rely is always true under the history ($\Sigma \vdash e_{\text{rely}} \Downarrow^\square \top$),
 283 then the body always satisfies the guarantee ($\Sigma \vdash e_{\text{guar}}[x := e_{\text{body}}] \Downarrow^\square \top$). Additionally, we
 284 can also assume that the valid properties in all three components hold, and we must also
 285 show that the unknown properties are valid. The fact that the checked semantics refers to
 286 a particular Σ is significant here: it allows the proof of contract validity to only consider

$$\boxed{\Sigma \vdash_{\pi} e \text{ valid}}$$

$$\begin{array}{c}
\overline{\Sigma \vdash_{\pi} v \text{ valid}} \text{ (CHKVALUE)} \quad \overline{\Sigma \vdash_{\pi} x \text{ valid}} \text{ (CHKVAR)} \\
\frac{\Sigma \vdash_{\pi} e_1 \text{ valid} \quad \dots \quad \Sigma \vdash_{\pi} e_n \text{ valid}}{\Sigma \vdash_{\pi} p(\bar{e}) \text{ valid}} \text{ (CHKPRIM)} \\
\frac{}{\sigma \vdash_{\pi} v \text{ fby } e' \text{ valid}} \text{ (CHKFBY}_1\text{)} \quad \frac{\text{length}(\Sigma) > 0 \quad \Sigma \vdash_{\pi} e' \text{ valid}}{\Sigma; \sigma \vdash_{\pi} v \text{ fby } e' \text{ valid}} \text{ (CHKFBY}_S\text{)} \\
\frac{\Sigma \vdash_{\pi} e[x := \mathbf{rec } x. e] \text{ valid}}{\Sigma \vdash_{\pi} \mathbf{rec } x. e[x] \text{ valid}} \text{ (CHKREC)} \quad \frac{\Sigma \vdash_{\pi} e'[x := e] \text{ valid}}{\Sigma \vdash_{\pi} \mathbf{let } x = e \text{ in } e'[x] \text{ valid}} \text{ (CHKLET)} \\
\frac{(\pi = \pi' \implies \Sigma \vdash e \Downarrow^{\square} \top) \quad \Sigma \vdash_{\pi} e \text{ valid}}{\Sigma \vdash_{\pi} \mathbf{check}_{\pi'} e \text{ valid}} \text{ (CHKCHECK)} \\
\frac{(\pi = \pi' \implies \Sigma \vdash e_{\text{rely}} \Downarrow^{\square} \top) \quad (\pi = \boxtimes \implies \Sigma \vdash e_{\text{rely}} \Downarrow^{\square} \top \implies \Sigma \vdash e_{\text{guar}}[x := e_{\text{body}}] \Downarrow^{\square} \top) \quad \Sigma \vdash_{\pi} e_{\text{rely}} \text{ valid}}{(\Sigma \vdash e_{\text{rely}} \Downarrow^{\square} \top \implies \Sigma \vdash_{\pi} e_{\text{body}} \text{ valid} \wedge \Sigma \vdash_{\pi} e_{\text{guar}}[x := e_{\text{body}}] \text{ valid})} \text{ (CHKCONTRACT)} \\
\frac{}{\Sigma \vdash_{\pi} \mathbf{contract}_{\pi'} \{e_{\text{rely}}\} e_{\text{body}} \{x. e_{\text{guar}}[x]\} \text{ valid}}
\end{array}$$

■ **Figure 5** Checked semantics for Pipit; the judgment form $\Sigma \vdash_{\pi} e \text{ valid}$ denotes that evaluating expression e under streaming history Σ satisfies the checks and rely-guarantee contract requirements that are labelled with property status π .

287 streaming histories where the rely actually holds. *TODO: WAFFLE? This is expanded on*
 288 *in Subsubsection 3.2.1; kill here?*

289 To prove that a contract *instantiation* (a call-site) is valid, we show that, under the calling
 290 environment, the rely clause is always true. Crucially, the proof can also use the fact that, if
 291 the rely is always true, then the guarantee is always true. This sort of feedback is necessary
 292 for proving properties of mutually-dependent calls. Although this feedback appears circular,
 293 we enforce causality by requiring that occurrences of recursive streams are guarded by delays
 294 (Subsection 3.3).

295 We define the checked semantics of Pipit in Figure 5. The checked semantics mostly
 296 follows the structure of the dynamic semantics, checking any properties and contracts as
 297 they are encountered.

298 Rules `CHKVALUE` and `CHKVAR` state that values and variables are always valid.

299 Rule `CHKPRIM` checks a primitive application by descending into the subexpressions.

300 Rules `CHKFBY1` and `CHKFBYS` are derived from the structure of the big-step rules `FBY1`
 301 and `FBYS`. At an input stream of length one, `CHKFBY1` asserts that all subproperties
 302 hold for the (non-existent) previous values in the stream. At subsequent parts of the
 303 stream, `CHKFBYS` discards the most recent element of the stream history and checks the
 304 subexpression with the previous inputs.

305 Rules `CHKREC` and `CHKLET` both perform the same unfolding as the corresponding
 306 big-step rules and check the resulting expression.

307 Finally, the heavy lifting is performed by rules `CHKCHECK` and `CHKCONTRACT`.

308 Rule `CHKCHECK` applies when checking property status π of an expression `check π' e`. If
 309 the check-expression has the same status as what we are checking ($\pi = \pi'$), then we perform
 310 the actual check by evaluating the expression e and requiring it to evaluate to a stream
 311 of trues. Otherwise, we do not need to evaluate the check-expression. In both cases, we
 312 descend into the expression and check its subexpressions, as they may have nested properties.
 313 Such nested properties are unlikely to be written directly by the user, but might occur after
 314 program transformations such as inlining.

315 Rule `CHKCONTRACT` applies when checking property status π of a contract with expression
 316 `contract π' {erely} ebody {x. eguar[x]}`. Although we only include one property status on
 317 the contract, conceptually there are two distinct properties: one for the caller (π') and one
 318 for the definition itself (assumed to be \square). To check the caller property when $\pi = \pi'$, we
 319 evaluate the rely e_{rely} and require it to be true. To check the definition property when $\pi = \square$,
 320 we assume that the rely holds, and check that the body satisfies the guarantee. We also
 321 descend into the subexpressions to check them; when checking the body and guarantee, we
 322 can assume that the rely holds. Unfortunately, this rule must deal with the two different
 323 roles of a contract at once; in the next section, we will separate the two roles.

324 3.2.1 Blessing expressions and contracts

325 Blessing is a meta-operation that replaces the property statuses in an expression so that all
 326 checks and contracts are marked as valid (\square). Blessing an expression requires a proof that
 327 the checked semantics hold for all input streams:

$$\frac{\forall \Sigma. \Sigma \vdash_{\square} e \text{ valid} \implies \Sigma \vdash_{\square} e \text{ valid}}{\text{bless } e} \text{ (BLESEXPRESSION)}$$

328 Blessing is slightly different for contract definitions, as we need to separate the definition
 329 of the contract from the instantiation. To check that a contract definition is valid, we show
 330 that if the rely clause is always true for a particular input, then the body satisfies the

331 guarantee for the same inputs. We also assume that the valid properties in the rely, body
 332 and guarantee hold, and show the corresponding unknown properties:

$$\begin{aligned} \text{let } \text{contract_valid } \{e_{\text{rely}}\} e_{\text{body}} \{e_{\text{guar}}\} : \text{prop} = \\ \forall \Sigma. (\Sigma \vdash_{\square} (e_{\text{rely}}, e_{\text{body}}, e_{\text{guar}}[x := e_{\text{body}}]) \text{ valid} \wedge \Sigma \vdash e_{\text{rely}} \Downarrow^{\square} \top) \\ \implies (\Sigma \vdash_{\square} (e_{\text{rely}}, e_{\text{body}}, e_{\text{guar}}[x := e_{\text{body}}]) \text{ valid} \wedge \Sigma \vdash e_{\text{guar}}[x := e_{\text{body}}] \Downarrow^{\square} \top) \end{aligned}$$

After proving that the contract is valid for all inputs, we can bless the contract definition. Blessing the contract definition blesses the subexpressions for the rely, body and guarantee, but leaves the contract's *instantiation* property status as unknown:

$$\frac{\text{contract_valid } \{e_{\text{rely}}\} e_{\text{body}} \{e_{\text{guar}}\}}{\text{bless_contract } \{e_{\text{rely}}\} e_{\text{body}} \{e_{\text{guar}}\}} (\text{BLESSCONTRACT})$$

333 3.3 Causality and metatheory

334 To ensure that recursive streams have a computational interpretation, we require that all
 335 recursive streams are guarded by a followed-by delay. We implement this as a simple syntactic
 336 check: each `rec x. e` can only mention `x` inside a followed-by. This check is stricter than
 337 necessary: for example, the expression `rec x. (let x' = x + 1 in 0 fby x')` does mention
 338 the recursive stream `x` outside of the delay, but after inlining the let, it would be causal. We
 339 hope to relax this restriction somewhat in future work.

340 The causality restriction gives us some important properties about the metatheory. The
 341 most important property is that the dynamic semantics form a total function: given a
 342 streaming history and a causal expression, we can evaluate the expression to a value. These
 343 properties are mechanised in F^* .

344 ► **Theorem 1 (bigstep-is-total).** *For any non-empty streaming history Σ and causal expression*
 345 *e , there exists some value v such that e evaluates to v ($\Sigma \vdash e \Downarrow v$).*

346 The relationship between substitution and the streaming history is also important. In
 347 general, we have a substitution property that states that evaluating a substituted expression
 348 $e[x := e']$ under some context Σ is equivalent to evaluating e' and adding it to the context Σ :

349 ► **Theorem 2 (bigstep-substitute).** *For a streaming history Σ and a causal expression e , if*
 350 *$e[x := e']$ evaluates to a value v ($\Sigma \vdash e \Downarrow v$), then we can evaluate e' to some stream V and*
 351 *extend the streaming history to evaluate e to the original value ($\Sigma[x \mapsto V] \vdash e \Downarrow v$). The*
 352 *converse is also true.*

353 The semantics in Figure 4 for a recursive expression `rec x. e` performs one step of
 354 recursion by substituting `x` for the recursive expression. An alternative semantics would be
 355 to have the environment outside the semantics invent a stream V such that if we extend the
 356 streaming history with $x \mapsto V$, then e evaluates to V itself. The above substitution theorem
 357 can be used to show that these two semantics are equivalent. Thanks to causality, we can
 358 additionally show that, when evaluating e with $x \mapsto V$, the most recent value in V does not
 359 affect the result. This fact can be used to “seed” evaluation by starting with an arbitrary
 360 value:

361 ► **Theorem 3 (bigstep-rec-causal).** *For a streaming history $\Sigma; \sigma$ and a causal recursive*
 362 *expression `rec x. e`, if ($\Sigma; \sigma \vdash e \Downarrow v$), then updating $\sigma[x]$ with any value v' results in the*
 363 *same value: ($\Sigma; \sigma[x \mapsto v'] \vdash e \Downarrow v$).*

```

type system (input:  $\Gamma$ ) (result:  $\tau$ ) = {
  state:  $\Gamma$ ;
  free:  $\Gamma$ ;
  init: heap state;
  step: heap input  $\rightarrow$  heap free  $\rightarrow$  heap state  $\rightarrow$  step_result state result;
}

type step_result (state:  $\Gamma$ ) (result:  $\tau$ ) = {
  update: heap state;
  value: result;
  rely: prop;
  guar: prop;
}

```

■ **Figure 6** Abstract transition system type definitions

364 4 Abstract transition systems

365 To prove properties about Pipit programs, we translate to an *abstract* transition system,
 366 so-called because it abstracts away the implementation details of contract instantiations. For
 367 extraction we also translate to *executable* transition systems, which we discuss in Section 5.

368 Figure 6 shows the types of transition systems. A transition system is parameterised by
 369 its input context and the result type. It also contains two internal contexts: firstly, the state
 370 context describes the private state required to execute the machine; secondly, the free context
 371 contains any extra input values that the transition system would like to quantify over. The
 372 free context is used to allow the system to ask for arbitrary values from the environment,
 373 when it would not otherwise be able to return a concrete value.

374 For contract instantiations, which abstract over the implementation, the natural transla-
 375 tion to a transition system would involve an existential quantifier: there exists some value
 376 that satisfies the specification. Unfortunately, such an existential quantifier requires a step
 377 *relation* rather than a step *function*. Using a step relation complicates the resulting transition
 378 system, as other operations such as primitive application must also introduce existential
 379 quantifiers; such quantifiers block normalisation and result in a more complex transition
 380 system. Instead, the free context provides the step function with a fresh unconstrained value
 381 of the desired type, which the step function can then constrain.

382 As usual, the step-result contains the updated state for the transition system, as well
 383 as the result value. The step-result additionally contains two propositions for the ‘rely’,
 384 or assumptions about the execution environment, and ‘guarantee’, or obligations that the
 385 transition system must show. For the transition system corresponding to an expression e ,
 386 these propositions are analogous to the known checked semantics $\Sigma \vdash_{\square} e$ valid and unknown
 387 checks $\Sigma \vdash_{\square} e$ valid respectively.

388 Our implementation includes a mechanised proof that, for causal expressions, the transition
 389 system is an abstraction of the original expression’s dynamic semantics. The proof that the
 390 rely and guarantee propositions correspond to the checked semantics is future work.

391 Figure 7 defines the internal state and free contexts required for an expression. For most
 392 expression forms, the state and free contexts are defined by taking the union of the contexts
 393 of subexpressions. Followed-by delays introduce a local state variable x_{fby} in which to store

$$\begin{aligned}
\llbracket v \rrbracket_{\text{state}} &= \cdot \\
\llbracket x \rrbracket_{\text{state}} &= \cdot \\
\llbracket p(\bar{e}) \rrbracket_{\text{state}} &= \bigcup_i \llbracket e_i \rrbracket_{\text{state}} \\
\llbracket v \text{ fby } e \rrbracket_{\text{state}} &= x_{\text{fby}} : \tau, \llbracket e \rrbracket_{\text{state}} && \text{(fresh } x_{\text{fby}}) \\
\llbracket \text{rec } x. e \rrbracket_{\text{state}} &= \llbracket e \rrbracket_{\text{state}} \\
\llbracket \text{let } x = e \text{ in } e' \rrbracket_{\text{state}} &= \llbracket e \rrbracket_{\text{state}} \cup \llbracket e' \rrbracket_{\text{state}} \\
\llbracket \text{check}_\pi e \rrbracket_{\text{state}} &= \llbracket e \rrbracket_{\text{state}} \\
\llbracket \text{contract}_\pi \{e_r\} e_b \{x. e_g\} \rrbracket_{\text{state}} &= \llbracket e_r \rrbracket_{\text{state}} \cup \llbracket e_b \rrbracket_{\text{state}} \\
\\
\llbracket v \rrbracket_{\text{free}} &= \cdot \\
\llbracket x \rrbracket_{\text{free}} &= \cdot \\
\llbracket p(\bar{e}) \rrbracket_{\text{free}} &= \bigcup_i \llbracket e_i \rrbracket_{\text{free}} \\
\llbracket v \text{ fby } e \rrbracket_{\text{free}} &= \llbracket e \rrbracket_{\text{free}} \\
\llbracket \text{rec } x. e \rrbracket_{\text{free}} &= x : \tau, \llbracket e \rrbracket_{\text{free}} \\
\llbracket \text{let } x = e \text{ in } e' \rrbracket_{\text{free}} &= \llbracket e \rrbracket_{\text{free}} \cup \llbracket e' \rrbracket_{\text{state}} \\
\llbracket \text{check}_\pi e \rrbracket_{\text{free}} &= \llbracket e \rrbracket_{\text{free}} \\
\llbracket \text{contract}_\pi \{e_r\} e_b \{x. e_g\} \rrbracket_{\text{free}} &= x : \tau, \llbracket e_r \rrbracket_{\text{free}} \cup \llbracket e_b \rrbracket_{\text{state}}
\end{aligned}$$

■ **Figure 7** Transition system typing contexts of expressions; for an expression e , $\llbracket e \rrbracket_{\text{state}} : \Gamma$ and $\llbracket e \rrbracket_{\text{free}} : \Gamma$ describe the heaps used to store the expression's internal state and extra inputs.

394 the most recent stream value. We generate a fresh variable here, though the implementation
395 uses de Bruijn indices. Recursive streams and contracts both introduce new bindings into
396 the free context, assuming that their binders x are unique.

397 Figure 8 defines the translation for expressions. Values and expressions have no internal
398 state. For variables, we look for the variable binding in either of the input or free heaps;
399 bindings are unique and cannot occur in both. We omit the rely and guarantee definitions
400 here; both are trivially true.

401 To translate primitives, we union together the initial states of the subexpressions; updating
402 the state is similar. For the rely and guarantee definitions, we take the conjunction: we can
403 assume that all subexpressions rely clauses hold, and must show that all guarantees hold.

404 To translate a followed-by $v \text{ fby } e$, we initialise the follow-by's unique binder x_{fby} to v .
405 At each step, we return the value in the local state, before updating the local state to the
406 subexpression's new value. The rely and guarantee differ from the checked semantics here:
407 in the checked semantics, we check the subexpression on the previous inputs, but here we
408 check the current subexpression. This means that a single step of the rely and guarantee do
409 not exactly correspond to the checked semantics; however, we posit that they are equivalent
410 for a rely and guarantee that has been proven to hold for any sequence of inputs.

411 To translate a recursive expression $\text{rec } x. e$ of type τ , we require an arbitrary value
412 $x : \tau$ in the free heap. The rely proposition constrains the free variable x to be the result of
413 evaluating e with the binding for x passed along, thus closing the recursive loop.

414 To translate let-expressions $\text{let } x = e \text{ in } e'$, we extend the input heap with the value of
415 e before evaluating e' . The presentation here duplicates the computation of the value of e ,
416 but this is not an issue in practice.

417 To translate a check property, we inspect the property status. If the property is known to
418 be valid, then we can assume the property is true in the rely clause. Otherwise, we include
419 the property as an obligation in the guarantee clause. In either case, we also include the
420 subexpression's rely and guarantee clauses.

421 Finally, to translate contract instantiations, we use the contract's rely and guarantee and

$$\begin{aligned}
\llbracket v \rrbracket_{\text{init}} &= () \\
\llbracket v \rrbracket_{\text{value}}(i, f, s) &= v \\
\\
\llbracket x \rrbracket_{\text{init}} &= () \\
\llbracket x \rrbracket_{\text{value}}(i, f, s) &= (i \cup f).x \\
\\
\llbracket p(\bar{e}) \rrbracket_{\text{init}} &= \bigcup_i \llbracket e_i \rrbracket_{\text{init}} \\
\llbracket p(\bar{e}) \rrbracket_{\text{value}}(i, f, s) &= \text{prim-sem}(p, \overline{\llbracket e \rrbracket_{\text{value}}(i, f, s)}} \\
\llbracket p(\bar{e}) \rrbracket_{\text{update}}(i, f, s) &= \bigcup_i \llbracket e_i \rrbracket_{\text{update}}(i, f, s) \\
\llbracket p(\bar{e}) \rrbracket_{\text{rely}}(i, f, s) &= \bigwedge_i \llbracket e_i \rrbracket_{\text{rely}}(i, f, s) \\
\llbracket p(\bar{e}) \rrbracket_{\text{guar}}(i, f, s) &= \bigwedge_i \llbracket e_i \rrbracket_{\text{guar}}(i, f, s) \\
\\
\llbracket v \text{ fby } e \rrbracket_{\text{init}} &= \llbracket e \rrbracket_{\text{init}} \cup \{x \text{ fby } \mapsto v\} \\
\llbracket v \text{ fby } e \rrbracket_{\text{value}}(i, f, s) &= s.x \text{ fby} \\
\llbracket v \text{ fby } e \rrbracket_{\text{update}}(i, f, s) &= \llbracket e \rrbracket_{\text{update}}(i, f, s) \cup \{x \text{ fby } \mapsto \llbracket e \rrbracket_{\text{value}}(i, f, s)\} \\
\llbracket v \text{ fby } e \rrbracket_{\text{rely}}(i, f, s) &= \llbracket e \rrbracket_{\text{rely}}(i, f, s) \\
\llbracket v \text{ fby } e \rrbracket_{\text{guar}}(i, f, s) &= \llbracket e \rrbracket_{\text{guar}}(i, f, s) \\
\\
\llbracket \text{rec } x. e \rrbracket_{\text{init}} &= \llbracket e \rrbracket_{\text{init}} \\
\llbracket \text{rec } x. e \rrbracket_{\text{value}}(i, f, s) &= f.x \\
\llbracket \text{rec } x. e \rrbracket_{\text{update}}(i, f, s) &= \llbracket e \rrbracket_{\text{update}}(i, f, s) \\
\llbracket \text{rec } x. e \rrbracket_{\text{rely}}(i, f, s) &= \llbracket e \rrbracket_{\text{rely}}(i, f, s) \\
&\wedge f.x = \llbracket e \rrbracket_{\text{value}}(i, f, s) \\
\llbracket \text{rec } x. e \rrbracket_{\text{guar}}(i, f, s) &= \llbracket e \rrbracket_{\text{guar}}(i, f, s) \\
\\
\llbracket \text{let } x = e \text{ in } e' \rrbracket_{\text{init}} &= \llbracket e \rrbracket_{\text{init}} \cup \llbracket e' \rrbracket_{\text{init}} \\
\llbracket \text{let } x = e \text{ in } e' \rrbracket_{\text{value}}(i, f, s) &= \llbracket e' \rrbracket_{\text{value}}(i \cup \{x \mapsto \llbracket e \rrbracket_{\text{value}}(i, f, s)\}, f, s) \\
\llbracket \text{let } x = e \text{ in } e' \rrbracket_{\text{update}}(i, f, s) &= \llbracket e' \rrbracket_{\text{update}}(i \cup \{x \mapsto \llbracket e \rrbracket_{\text{value}}(i, f, s)\}, f, s) \\
&\cup \llbracket e \rrbracket_{\text{update}}(i, f, s) \\
\llbracket \text{let } x = e \text{ in } e' \rrbracket_{\text{rely}}(i, f, s) &= \llbracket e' \rrbracket_{\text{rely}}(i \cup \{x \mapsto \llbracket e \rrbracket_{\text{value}}(i, f, s)\}, f, s) \\
&\wedge \llbracket e \rrbracket_{\text{rely}}(i, f, s) \\
\llbracket \text{let } x = e \text{ in } e' \rrbracket_{\text{guar}}(i, f, s) &= \llbracket e' \rrbracket_{\text{guar}}(i \cup \{x \mapsto \llbracket e \rrbracket_{\text{value}}(i, f, s)\}, f, s) \\
&\wedge \llbracket e \rrbracket_{\text{guar}}(i, f, s) \\
\\
\llbracket \text{check}_\pi e \rrbracket_{\text{init}} &= \llbracket e \rrbracket_{\text{init}} \\
\llbracket \text{check}_\pi e \rrbracket_{\text{value}}(i, f, s) &= () \\
\llbracket \text{check}_\pi e \rrbracket_{\text{update}}(i, f, s) &= \llbracket e \rrbracket_{\text{update}}(i, f, s) \\
\llbracket \text{check}_\pi e \rrbracket_{\text{rely}}(i, f, s) &= (\pi = \square \implies \llbracket e \rrbracket_{\text{value}}(i, f, s)) \wedge \llbracket e \rrbracket_{\text{rely}}(i, f, s) \\
\llbracket \text{check}_\pi e \rrbracket_{\text{guar}}(i, f, s) &= (\pi = \boxtimes \implies \llbracket e \rrbracket_{\text{value}}(i, f, s)) \wedge \llbracket e \rrbracket_{\text{guar}}(i, f, s) \\
\\
\llbracket \text{contract}_\pi \{e_r\} e_b \{x. e_g\} \rrbracket_{\text{init}} &= \llbracket e_r \rrbracket_{\text{init}} \cup \llbracket e_g \rrbracket_{\text{init}} \\
\llbracket \text{contract}_\pi \{e_r\} e_b \{x. e_g\} \rrbracket_{\text{value}}(i, f, s) &= f.x \\
\llbracket \text{contract}_\pi \{e_r\} e_b \{x. e_g\} \rrbracket_{\text{update}}(i, f, s) &= \llbracket e_r \rrbracket_{\text{update}}(i, f, s) \cup \llbracket e_g \rrbracket_{\text{update}}(i, f, s) \\
\llbracket \text{contract}_\pi \{e_r\} e_b \{x. e_g\} \rrbracket_{\text{rely}}(i, f, s) &= (\llbracket e_r \rrbracket_{\text{value}}(i, f, s) \implies \llbracket e_g \rrbracket_{\text{value}}(i, f, s)) \\
&\wedge (\pi = \square \implies \llbracket e_r \rrbracket_{\text{value}}(i, f, s)) \\
&\wedge \llbracket e_r \rrbracket_{\text{rely}}(i, f, s) \wedge \llbracket e_g \rrbracket_{\text{rely}}(i, f, s) \\
\llbracket \text{contract}_\pi \{e_r\} e_b \{x. e_g\} \rrbracket_{\text{guar}}(i, f, s) &= (\pi = \boxtimes \implies \llbracket e_r \rrbracket_{\text{value}}(i, f, s)) \\
&\wedge \llbracket e_r \rrbracket_{\text{guar}}(i, f, s) \wedge \llbracket e_g \rrbracket_{\text{guar}}(i, f, s)
\end{aligned}$$

■ **Figure 8** Transition system semantics; for an expression $\Gamma \vdash e : \tau$, $\llbracket e \rrbracket_{\text{init}} : \text{heap}$ $\llbracket e \rrbracket_{\text{state}}$ is the initial state. For each field of the step-result type, we define a translation function that takes the input, free and state heaps: for example, we define the value-result of a step with type $\llbracket e \rrbracket_{\text{value}} : \text{heap } \Gamma \rightarrow \text{heap } \llbracket e \rrbracket_{\text{free}} \rightarrow \text{heap } \llbracket e \rrbracket_{\text{state}} \rightarrow \tau$.

$$\boxed{\Sigma \vdash e \sim s}$$

$$\begin{array}{c}
\frac{}{\Sigma \vdash v \sim s} \text{ (IVALUE)} \qquad \frac{}{\Sigma \vdash x \sim s} \text{ (IVAR)} \\
\frac{\Sigma \vdash e_1 \sim s \quad \dots \quad \Sigma \vdash e_n \sim s}{\Sigma \vdash p(\bar{e}) \sim s} \text{ (IPRIM)} \qquad \frac{s.x\mathbf{fby} = v \quad \cdot \vdash e' \sim s}{\cdot \vdash v \mathbf{fby} e' \sim s} \text{ (IFBY}_0\text{)} \\
\frac{\Sigma; \sigma \vdash e' \Downarrow s.x\mathbf{fby} \quad \Sigma; \sigma \vdash e' \sim s}{\Sigma; \sigma \vdash v \mathbf{fby} e' \sim s} \text{ (IFBY}_S\text{)} \\
\frac{\Sigma \vdash \mathbf{rec} \ x. e \Downarrow^* V \quad \Sigma[x \mapsto V] \vdash e \sim s}{\Sigma \vdash \mathbf{rec} \ x. e[x] \sim s} \text{ (IREC)} \\
\frac{\Sigma \vdash e \Downarrow^* V \quad \Sigma \vdash e \sim s \quad \Sigma[x \mapsto V] \vdash e' \sim s}{\Sigma \vdash \mathbf{let} \ x = e \ \mathbf{in} \ e'[x] \sim s} \text{ (ILET)} \\
\frac{\Sigma \vdash e \sim s}{\Sigma \vdash \mathbf{check}_\pi e \sim s} \text{ (ICHECK)} \\
\frac{\Sigma \vdash e_{\mathbf{body}} \Downarrow^* V \quad \Sigma \vdash e_{\mathbf{rely}} \sim s \quad \Sigma[x \mapsto V] \vdash e_{\mathbf{guar}} \sim s}{\Sigma \vdash \mathbf{contract}_\pi \{e_{\mathbf{rely}}\} e_{\mathbf{body}} \{x. e_{\mathbf{guar}}[x]\} \sim s} \text{ (ICONTRACT)}
\end{array}$$

■ **Figure 9** Transition system state invariant

422 ignore the body. As with recursive expressions, we require an arbitrary value $x : \tau$ in the
423 free heap. The translation's rely allows us to assume that the contract definition holds: that
424 is, the contract's rely implies the contract's guarantee. If the contract instantiation is known
425 to be valid, we can also assume that the contract's rely holds. Otherwise, we include the
426 contract's rely as an obligation by putting it in the translation's guarantee.

427 In the contract instantiation, we assume that if the contract rely is true *at the current*
428 *step*, then the contract guarantee also holds at the current step. The true semantics of
429 the contract, however, only holds if the contract rely is true *at every step so far*. This
430 simplification is justified as our definition of validity for a transition system also requires
431 the translation rely to be true at every step. This simplification is essentially an application
432 of the $\Box\Box p \implies \Box p$ axiom of modal logic. The mechanised proof of this simplification is
433 future work.

434 **4.1 Translation correctness proofs**

435 We prove that the transition system is an abstraction of the dynamic semantics: that is, if
436 the expression evaluates to v under some context, then there exists some execution of the
437 transition system that also results in v . The transition system itself is deterministic, but the
438 free context provides the non-determinism; our theorem statement existentially quantifies
439 the free heap.

440 The results presented here rely heavily on the totality and substitution metaproperties
441 described in Subsection 3.3. Figure 9 defines the invariant for the abstraction proof; the
442 judgment form $\Sigma \vdash e \sim s$ checks that s is a valid state heap. We use the invariant to state

443 that, if executing the transition system for e on the entire streaming history Σ results in
 444 state heap s , then s is a valid state.

445 As most expressions do not modify the state heap, the invariant for most expressions
 446 simply descends into the subexpressions. Where new bindings are added, we use the dynamic
 447 semantics to extend the context with the new values. The invariant for follow-by expressions
 448 asserts that the initial state of the follow-by is the default value; on subsequent steps, the
 449 state corresponds to the dynamic semantics.

450 ► **Theorem 4** (translation-abstraction). *For a well-typed causal expression e and streaming*
 451 *history Σ , if e evaluates to v ($\Sigma \vdash e \Downarrow v$), then there exists a sequence of free heaps Σ_F such*
 452 *that repeated application of the transition system’s step results in v .*

453 5 Extraction

454 Pipit can generate executable code which is suitable for real-time execution on embedded
 455 devices. The code extraction uses a variation of the abstract transition system described in
 456 Section 4, with two main differences to ensure that the result is executable without relying
 457 on the environment to provide values for the free context. Contracts are straightforward to
 458 execute by using the body of the contract rather than abstracting over the implementation.

459 To execute recursive expressions $\text{rec } x. e : \tau$, we require an arbitrary value of type τ to
 460 seed the fixpoint, as described in Subsection 3.3. We first call the step function to evaluate e
 461 with x bound to \perp_τ . This step call returns the correct value, but the updated state is invalid,
 462 as it may refer to the bottom value. To get the correct state, we call the step function again,
 463 this time with e bound to v .

464 This translation to transition systems is verified to preserve the original semantics. To
 465 extract the program, we use a *hybrid embedding* as described in [23], which is similar to staged-
 466 compilation. The hybrid embedding involves a deep embedding of the Pipit core language,
 467 while the translation to executable transition systems produces a shallow embedding. We
 468 use the F^* host language’s normalisation-by-evaluation and tactic support [31] to specialise
 469 the application of the translation to a particular input program. This specialisation results
 470 in a concrete transition system that fits in the Low^* [34] subset of F^* , which can then be
 471 extracted to statically-allocated C code.

472 The translation for recursive streams described above calls the step function of the sub-
 473 stream twice, which can duplicate work. The normalisation strategy used to partially-evaluate
 474 the translation inlines the two occurrences of the step function, and is often able to remove the
 475 duplicate work, but this removal is not guaranteed. Our current approach is also unsuitable
 476 for generating imperative array code, as our shallowly-embedded pure transition system
 477 requires pure arrays. In the future, we intend to address array computations and the above
 478 work duplication by introducing an intermediate imperative language such as `Obc` [3], a static
 479 object-based language suitable for synchronous systems. Even with an added intermediate
 480 language, we believe that a variant of our current translation and proof-of-correctness will
 481 remain useful as an intermediate semantics.

482 6 Evaluation

483 As a preliminary evaluation of Pipit, we have implemented the high-level logic of a time-
 484 triggered Controller Area Network (CAN) bus driver [1]. The CAN bus is commonly found
 485 in safety-critical automotive and industrial settings. The time-triggered network architecture

486 defines a static schedule of network traffic. All nodes on the network must adhere to the
 487 same schedule, which significantly increases the reliability of periodic messages [15].

488 At a high level, the schedule is described by a system matrix which consists of rows of
 489 basic cycles. Each basic cycle consists of a sequence of actions to be performed at particular
 490 time-marks. Actions in the schedule may not be relevant to all nodes, so each node has its
 491 own local array containing the relevant triggers; trigger actions include sending and receiving
 492 application-specific messages, sending reference messages, and triggering ‘watch’ alerts. The
 493 trigger action for receiving an application-specific message checks that a particular message
 494 has been received since the trigger was last executed; depending on this, the driver increments
 495 or decrements a message-status-counter, which will in turn signal an error once the upper
 496 limit is reached. Reference messages start a new basic cycle and are used to synchronise the
 497 nodes. Watch alerts are generally placed after the expected end of the cycle and are used to
 498 signal an error if no reference message is received.

499 The TTCAN protocol can be implemented in two levels of increasing complexity. In the
 500 first level, reference messages contain the index of the newly-started cycle. In the second
 501 level, the reference messages also contain the value of a global fractional clock and whether
 502 any gaps have occurred in the global clock, which allows other nodes to calibrate their own
 503 clocks. We implement the first level as it is more amenable to software implementation [22].

504 The implementation defines a streaming function that takes a stream describing the current
 505 time, the state of the hardware, and any received messages. It returns a stream of commands
 506 to be performed, such as sending a particular reference message. The implementation defines
 507 a pure streaming function. To actually interact with the hardware we assume a small
 508 hardware-interop layer that reads from the hardware registers and translates the commands
 509 to hardware-register writes, but we have not yet implemented this. We package the driver’s
 510 inputs into a record for convenience:

```

type driver_input = {
  local_time: network_time_unit;
  mode_cmd: option mode;
  tx_status: tx_status;
  bus_status: bus_status;
  rx_ref:    option ref_message;
  rx_app:    option app_message_index;
}

```

511 Here, the local-time field denotes the time-since-boot in *network time units*, which are
 512 based on the bitrate of the underlying network bus. The mode-command is an optional field
 513 which indicates requests from the application to enter configuration or execution mode. The
 514 transmission-status describes the status of the last transmission request and may be none,
 515 success, or various error conditions. The bus-status describes whether the bus is currently
 516 idle, busy, or in an error state. The two receive fields denote messages received from the bus;
 517 for application-specific messages the time-triggered logic only needs the message identifier.

518 The driver-logic returns a stream of commands for the hardware-interop layer to perform:

```

type commands = {
  enable_acks: bool;
  tx_ref:      option ref_message;
  tx_app:      option app_message_index;
  tx_delay:    network_time_unit;
}

```

519 The enable-acks field denotes whether the hardware should respond to messages from
520 other nodes with an acknowledgement bit; in the case of a severe error acknowledgements are
521 disabled, as the node must not write to the bus at all. The transmit fields denote whether
522 to send a reference message or an application-specific message. For application-specific
523 messages, the hardware-interop layer maintains the transmission buffers containing the actual
524 message payload. To meet the schedule as closely as possible, the driver anticipates the next
525 transmission and includes a transmission delay to tell the hardware exactly when to send the
526 next message.

527 6.1 Runtime

528 The implementation includes an extension of the trigger-fetch logic described in Section 2, as
529 well as state machines for tracking node synchronisation, master status and fault handling.
530 We generate real-time C code as described in Section 5. We evaluated the generated C code
531 by executing with randomised inputs and measuring the worst-case-execution-time on a
532 Raspberry Pi Pico (RP2040) microcontroller. The runtime of the driver logic is fairly stable:
533 over 5,000 executions, the measured worst-case execution time was $114\mu s$, while the average
534 was $107\mu s$ with a standard deviation of $2.3\mu s$. Earlier work on fault-tolerant TTCAN [41]
535 describes the required slot sizes — the minimum time between triggers — to achieve bus
536 utilisation at different bus rates. For a 125Kbit/s bus, a slot size of approximately $1,500\mu s$
537 is required to achieve utilisation above 85 per cent. For the maximum CAN bus rate of
538 1Mbit/s, the required slot size is $184\mu s$. Further evaluation is required to ensure that the
539 complete runtime including the hardware-interop layer is sufficient for full-speed CAN.

540 Our code generation can be improved in a few ways. A common optimisation in Lustre is
541 to fuse consecutive if-statements with the same condition [5]; such an optimisation seems
542 useful here, as our treatment of optional values introduces repeated unpacking and repacking.
543 Some form of array fusion [37] may also be useful for removing redundant array operations.
544 Our current extraction generates a transition-system with a step function which returns
545 a tuple of the updated state and result. Composing these step functions together results
546 in repeated boxing and unboxing of this tuple; we currently rely on the F^* normaliser to
547 remove this boxing. In the future, we plan to build on the current proofs to implement a
548 more-sophisticated encoding that introduces less overhead.

549 6.2 Verification

550 We have verified a simplified trigger-fetch mechanism, as presented earlier (Section 2). For
551 comparison, we implemented the same logic in the Kind2 model-checker [11]. The restrictions
552 placed on the triggers array — that triggers are sorted by time-mark, that there must be an
553 adequate time-gap between a trigger and its next-enabled, and that a trigger's time-mark
554 must be greater-than-or-equal-to its index — are naturally expressed with quantifiers. The
555 Kind2 model-checker includes experimental array and quantifier support [26]. Due to the
556 experimental nature of these features, we had to work around some limitations: for example,
557 the use of arrays and quantifiers disables IC3-based invariant generation; quantified variables
558 cannot be used in function calls; and the use of top-level constant arrays caused runtime
559 errors that rendered most properties invalid [27].

560 We were able to verify the Kind2 implementation of the simplified trigger-fetch mechanism
561 for trigger arrays containing up to 16 elements; above that, a 32-size array reported multiple
562 runtime errors and did not terminate after several hours. For reference, the `M_TTCAN`
563 hardware implementation of TTCAN supports up to 64 triggers [36].

| size | Kind2 | | | | Pipit | |
|------|-------------------|------|-----------------|-----|------------|-----|
| | simple enable-set | | full enable-set | | wall-clock | CPU |
| | wall-clock | CPU | wall-clock | CPU | | |
| 1 | 2s | 3s | 6s | 12s | 6s | 6s |
| 2 | 3s | 3s | 8s | 12s | 6s | 6s |
| 4 | 5s | 11s | 12s | 24s | 6s | 6s |
| 8 | 8s | 13s | 82s | 90s | 6s | 6s |
| 16 | 125s | 276s | error | | 6s | 6s |
| 32 | error | | error | | 6s | 6s |

■ **Figure 10** Verification time for trigger-fetch; simple enable-set uses a simplified version of the enable-set, while full enable-set uses bitwise arithmetic as in the TTCAN specification.

564 We made a critical simplification in the Kind2 implementation, which was to modify
 565 the trigger-enabled set to be a single cycle index. In the specification, the enabled set is
 566 implemented as a cycle-offset and repeat-factor. Checking if a trigger is enabled in the
 567 current cycle requires nonlinear arithmetic, which is difficult for SMT solvers. In our Pipit
 568 development, we can treat the definition of the cycle set abstractly. However, in the Kind2
 569 development, quantifiers cannot contain function calls, which means that we cannot hide the
 570 implementation of the enabled-set check by providing an abstract contract. This limitation
 571 also makes the specification quite unwieldy, as functions must be manually inlined.

572 Figure 10 shows the verification runtime for different sizes of arrays; the Pipit version
 573 is parametric in the array size, and is thus verified for all sizes of arrays. We ran these
 574 experiments on a 2020 M1 MacBook Air with 16 gigabytes of RAM. Both Kind2 and Pipit
 575 developments of the simplified trigger-fetch logic are roughly the same size, on the order of
 576 two-hundred lines of code including comments.

577 We plan to verify the remainder of the TTCAN implementation and publish it separately.
 578 Prior work formalising TTCAN has variously modeled the protocol itself [39, 33, 30], instances
 579 of the protocol [20], and abstract models of TTCAN implementations [29], but we are unaware
 580 of any prior work that has verified an *executable* implementation of TTCAN.

581 Separately, Pipit has also been used to implement and verify a real-time controller for a
 582 coffee machine reservoir control system [38]. The reservoir has a float switch to sense the
 583 water level and a solenoid to allow the intake of water. The specification includes a simple
 584 model of the water reservoir and shows that the reservoir does not exceed the maximum
 585 level under different failure-mode assumptions.

586 7 Related work

587 Using existing Lustre tools to verify *and* execute the time-triggered CAN driver from Section 2
 588 is nontrivial. Compiling the triggers array with an unverified compiler such as Lustre V6 [24]
 589 or Heptagon [19] is straightforward; however, the verified Lustre compiler Vélus [7] does not
 590 support arrays or a foreign-function interface. Recent work on translation validation for
 591 LustreC [9] also does not yet support arrays.

592 Verifying the time-triggered CAN driver is trickier, as the restrictions placed on the
 593 triggers array — that triggers are sorted by time-mark, there must be an adequate time-gap
 594 between a trigger and its next-enabled, and a trigger’s time-mark must be greater-than-or-
 595 equal-to its index — naturally require quantifiers. As described in Section 6, the Kind2 does
 596 include experimental array and quantifier support, but is limited to verifying small arrays

597 up to 8 or 16 triggers. Additionally, due to the limitations on top-level array definitions,
 598 compiling the program with Lustre V6 would result in multiple copies of the entire triggers
 599 array on the stack.

600 Other model-checkers for Lustre such as Lesar [35], JKind [16] and the original Kind [21]
 601 do not support quantifiers. It may be possible to encode the quantifiers as fixed-size loops,
 602 but ensuring that these loops do not affect the execution or runtime complexity of the
 603 generated code does not appear to be straightforward.

604 These model-checkers have definite usability advantages over the general-purpose-prover
 605 approach offered here: they can often generate concrete counterexamples and implement
 606 counterexample-based invariant-generation techniques such as ICE [18] and PDR [8, 14].
 607 However, even when the problem can be expressed, these model-checkers do not provide much
 608 assurance that the semantics they use for proofs matches the compiled code. In the future, we
 609 would like to investigate integrating Pipit with a model-checker via an unverified extraction:
 610 such an extraction may allow some of the usability benefits such as counterexamples and
 611 invariant generation. If this integration were used solely for debugging and suggesting
 612 candidate invariants, then such a change would not expand the trusted computing base.

613 Recent work has also introduced a form of refinement types for Lustre [12]. Rather
 614 than using transition systems, this work generates self-contained verification conditions
 615 based on the types of streams. Such a type-based approach promises to allow abstraction
 616 of the implementation details. However, for general-purpose functions such as *count_when*
 617 from Section 2, it is not clear how to give it a specification that actually *abstracts* the
 618 implementation: a simple specification that the result is within some range would hide be
 619 insufficient for verifying the rest of the system. For this function, the best specification is
 620 likely to include a re-statement of the implementation itself.

621 The embedded language Copilot generates real-time C code for runtime monitoring [28].
 622 Recent work has used translation validation to show that the generated C code matches
 623 the high-level semantics [40]. Copilot supports model-checking via Kind2; however, the
 624 model-checking has a limited specification language and does not support contracts.

625 Early work embedding a denotational semantics of Lucid Synchronic in an interactive
 626 theorem prover focussed on the semantics itself, rather than proving programs [4]. There is
 627 ongoing work to construct a denotational semantics of Vélus for program verification [6]. We
 628 believe that the hybrid SMT approach of F* will allow for a better mixture of automated
 629 proofs with manual proofs. Compared to Vélus alone, the trusted computing base of Pipit is
 630 larger: we depend on all of F*, Low*'s C code extraction and the Z3 SMT solver.

631 The deferred aspect of our proofs is similar to the deferred proofs of verification conditions
 632 for imperative programs, such as [32]. However, such verification conditions are *syntactically*
 633 deferred so that the verification condition can be proved later; in our case, the verification
 634 conditions are *semantically* deferred, so that more knowledge of the enclosing program
 635 can be exploited in the proof. In imperative programs, this sort of extra knowledge is
 636 generally provided explicitly as loop invariants, and non-looping statements have their
 637 weakest precondition computed automatically. In Lustre-style reactive languages such as
 638 ours, programs tend to be composed of many nested recursive streams, which perform a
 639 similar function to loops. Explicitly specifying an invariant for each recursive stream would
 640 be cumbersome; deferring the proof allows such invariants to be implicit.

641 **8 Conclusion**

642 *TODO: requires rewrite TODO: future work: clocks*

643 Our preliminary results show that F*'s proof automation and code extraction are suitable
 644 for verifying reactive systems and executing them in real-time; these results still require
 645 further work. Next, we intend to verify the imperative code generation. Finally, we need
 646 to evaluate Pipit on larger control systems before extending the language to support more
 647 features, such as Lustre's clocks for describing partially-defined streams [10].

648 We are interested in further pursuing the intersection of model-checking with interactive
 649 theorem proving. A smart contract called Djed [42] currently uses a mixture of Kind2 [11]
 650 and manual Isabelle/HOL proofs to show that the contract is well-behaved. In future work,
 651 we would like to further investigate whether Pipit's integration of streaming proofs with F*'s
 652 automated proof system would be able to provide similar proofs, without introducing any
 653 semantic gap between the two systems.

654 Our current array support is limited: constant arrays provide a pure index function,
 655 while we only support fixed-size mutable arrays by wrapping bit-vectors. Array support is
 656 an obvious direction for future work; integrating with a verified array-fusion system such as
 657 [37] would be an interesting and useful extension.

658 — References —

- 659 1 ISO/CD 11898-4. Road vehicles - Controller area network (CAN) - Part 4: Time triggered
 660 communication. Standard, International Organization for Standardization, 2000.
- 661 2 Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The Satisfiability Modulo Theories Library
 662 (SMT-LIB). www.SMT-LIB.org, 2016.
- 663 3 Dariusz Biernacki, Jean-Louis Colaço, Grégoire Hamon, and Marc Pouzet. Clock-directed
 664 modular code generation for synchronous data-flow languages. In *Proceedings of the 2008
 665 ACM SIGPLAN-SIGBED conference on Languages, compilers, and tools for embedded systems*,
 666 pages 121–130, 2008.
- 667 4 Sylvain Boulmé and Grégoire Hamon. A clocked denotational semantics for Lucid-Synchrone
 668 in Coq. *Rap. tech., LIP6*, 2001.
- 669 5 Timothy Bourke, Lélío Brun, Pierre-Évariste Dagand, Xavier Leroy, Marc Pouzet, and Lionel
 670 Rieg. A formally verified compiler for Lustre. In *Proceedings of the 38th ACM SIGPLAN
 671 Conference on Programming Language Design and Implementation*, 2017.
- 672 6 Timothy Bourke, Paul Jeanmaire, and Marc Pouzet. Towards a denotational semantics of
 673 streams for a verified Lustre compiler. 2022. URL: [https://types22.inria.fr/files/2022/
 674 06/TYPES_2022_slides_28.pdf](https://types22.inria.fr/files/2022/06/TYPES_2022_slides_28.pdf).
- 675 7 Timothy Bourke, Basile Pesin, and Marc Pouzet. Verified compilation of synchronous dataflow
 676 with state machines. *ACM Transactions on Embedded Computing Systems*, 22(5s):1–26, 2023.
- 677 8 Aaron R Bradley. SAT-based model checking without unrolling. In *Verification, Model
 678 Checking, and Abstract Interpretation: 12th International Conference, VMCAI 2011, Austin,
 679 TX, USA, January 23-25, 2011. Proceedings 12*. Springer, 2011.
- 680 9 Lélío Brun, Christophe Garion, Pierre-Loïc Garoche, and Xavier Thirioux. Equation-directed
 681 axiomatization of lustre semantics to enable optimized code validation. *ACM Transactions on
 682 Embedded Computing Systems*, 22(5s):1–24, 2023.
- 683 10 Paul Caspi and Marc Pouzet. A functional extension to Lustre. *Intensional Programming I*,
 684 1995.
- 685 11 Adrian Champion, Alain Melsout, Christoph Stickse, and Cesare Tinelli. The Kind 2 model
 686 checker. In *Computer Aided Verification*, 2016.
- 687 12 Jiawei Chen, José Luiz Vargas de Mendonça, Shayan Jalili, Bereket Ayele, Bereket Ngussie
 688 Bekele, Zhemín Qu, Pranjal Sharma, Tigist Shiferaw, Yicheng Zhang, and Jean-Baptiste
 689 Jeannin. Synchronous programming and refinement types in robotics: From verification to
 690 implementation. In *Proceedings of the 8th ACM SIGPLAN International Workshop on Formal
 691 Techniques for Safety-Critical Systems*, 2022.

- 692 13 Jean-Louis Colaço, Bruno Pagano, and Marc Pouzet. Scade 6: A formal language for embedded
693 critical software development. In *2017 International Symposium on Theoretical Aspects of*
694 *Software Engineering (TASE)*, pages 1–11. IEEE, 2017.
- 695 14 Niklas Eén, Alan Mishchenko, and Robert Brayton. Efficient implementation of property
696 directed reachability. In *2011 Formal Methods in Computer-Aided Design (FMCAD)*. IEEE,
697 2011.
- 698 15 Thomas Fuehrer, Bernd Mueller, Florian Hartwich, and Robert Hugel. Time triggered CAN
699 (TTCAN). *SAE transactions*, pages 143–149, 2001.
- 700 16 Andrew Gacek, John Backes, Mike Whalen, Lucas Wagner, and Elaheh Ghassabani. The
701 JKind model checker. In *Computer Aided Verification: 30th International Conference, CAV*
702 *2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17,*
703 *2018, Proceedings, Part II 30*, pages 20–27. Springer, 2018.
- 704 17 Emilio Jesús Gallego Arias, Pierre Jouvelot, Sylvain Ribstein, and Dorian Desblancs. The
705 W-calculus: a synchronous framework for the verified modelling of digital signal processing
706 algorithms. In *Proceedings of the 9th ACM SIGPLAN International Workshop on Functional*
707 *Art, Music, Modelling, and Design*, pages 35–46, 2021.
- 708 18 Pranav Garg, Christof Löding, Parthasarathy Madhusudan, and Daniel Neider. ICE: A
709 robust framework for learning invariants. In *Computer Aided Verification: 26th International*
710 *Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna,*
711 *Austria, July 18-22, 2014. Proceedings 26*. Springer, 2014.
- 712 19 Léonard Gérard, Adrien Guatto, Cédric Pasteur, and Marc Pouzet. A modular memory
713 optimization for synchronous data-flow languages: application to arrays in a Lustre compiler.
714 *ACM SIGPLAN Notices*, 47(5), 2012.
- 715 20 Xiaoyun Guo, Toshiaki Aoki, and Hsin-Hung Lin. Model checking of in-vehicle networking
716 systems with CAN and FlexRay. *Journal of Systems and Software*, 161:110461, 2020.
- 717 21 George Hagen and Cesare Tinelli. Scaling up the formal verification of Lustre programs with
718 SMT-based techniques. In *2008 Formal Methods in Computer-Aided Design*. IEEE, 2008.
- 719 22 Florian Hartwich, Thomas Führer, Bernd Müller, and Robert Hugel. Integration of time
720 triggered CAN (TTCAN_TC). *SAE Transactions*, pages 112–119, 2002.
- 721 23 Son Ho, Jonathan Protzenko, Abhishek Bichhawat, and Karthikeyan Bhargavan. Noise*: A
722 library of verified high-performance secure channel protocol implementations. In *2022 IEEE*
723 *Symposium on Security and Privacy (SP)*, pages 107–124. IEEE, 2022.
- 724 24 Erwan Jahier, Pascal Raymond, and Nicolas Halbwachs. The Lustre V6 reference manual.
725 *Verimag, Grenoble, Dec*, 2016.
- 726 25 Kind2. Integer division rounds to negative infinite. Github issues, 2023. URL: <https://github.com/kind2-mc/kind2/issues/978>.
- 727 728 26 Kind2. *Kind2 user documentation*, 2.1.1 edition, 2023. URL: https://kind.cs.uiowa.edu/kind2_user_doc/doc.pdf.
- 729 730 27 Kind2. Top-level array definition causes runtime failures. Github issues, 2024. URL: <https://github.com/kind2-mc/kind2/issues/1043>.
- 731 732 28 Jonathan Laurent, Alwyn Goodloe, and Lee Pike. Assuring the guardians. In *Runtime*
733 *Verification: 6th International Conference, RV 2015, Vienna, Austria, September 22-25, 2015.*
734 *Proceedings*. Springer, 2015.
- 735 736 29 Gabriel Leen and Donal Heffernan. Modeling and verification of a time-triggered networking
737 protocol. In *International Conference on Networking, International Conference on Systems*
738 *and International Conference on Mobile Communications and Learning Technologies (IC-*
739 *NICONSML'06)*, pages 178–178. IEEE, 2006.
- 740 741 30 Xin Li, Jian Guo, Yongxin Zhao, and Xiaoran Zhu. Formal modeling and verifying the
742 TTCAN protocol from a probabilistic perspective. *Journal of Circuits, Systems and Computers*,
743 28(10):1950177, 2018.
- 742 743 31 Guido Martínez, Danel Ahman, Victor Dumitrescu, Nick Giannarakis, Chris Hawblitzel,
Cătălin Hrițcu, Monal Narasimhamurthy, Zoe Paraskevopoulou, Clément Pit-Claudel, Jonathan

- 744 Protzenko, et al. Meta-F*: Proof automation with SMT, tactics, and metaprograms. In
745 *Programming Languages and Systems: 28th European Symposium on Programming, ESOP*
746 *2019, Held as Part of the European Joint Conferences on Theory and Practice of Software,*
747 *ETAPS 2019, Prague, Czech Republic, April 6–11, 2019, Proceedings.* Springer International
748 Publishing Cham, 2019.
- 749 **32** Liam O’Connor. Deferring the details and deriving programs. In *Proceedings of the 4th ACM*
750 *SIGPLAN International Workshop on Type-Driven Development*, pages 27–39, 2019.
- 751 **33** Can Pan, Jian Guo, Longfei Zhu, Jianqi Shi, Huibiao Zhu, and Xinyun Zhou. Modeling and
752 verification of CAN bus with application layer using UPPAAL. *Electronic Notes in Theoretical*
753 *Computer Science*, 309:31–49, 2014.
- 754 **34** Jonathan Protzenko, Jean Karim Zinzindohoué, Aseem Rastogi, Tahina Ramananandro, Peng
755 Wang, Santiago Zanella Béguelin, Antoine Delignat-Lavaud, Catalin Hritcu, Karthikeyan
756 Bhargavan, Cédric Fournet, et al. Verified low-level programming embedded in F*. *Proc.*
757 *ACM program. lang.*, 1(ICFP), 2017.
- 758 **35** Pascal Raymond. Synchronous program verification with Lustre/Lesar. *Modeling and Verific-*
759 *ation of Real-Time Systems*, 2008.
- 760 **36** Robert Bosch GmbH. *M_TTCAN Time-triggered Controller Area Network User’s Manual*,
761 3.3.0 edition, 2019. URL: [https://www.bosch-semiconductors.com/media/ip_modules/pdf_](https://www.bosch-semiconductors.com/media/ip_modules/pdf_2/m_can/mttcan_users_manual_v330.pdf)
762 [2/m_can/mttcan_users_manual_v330.pdf](https://www.bosch-semiconductors.com/media/ip_modules/pdf_2/m_can/mttcan_users_manual_v330.pdf).
- 763 **37** Amos Robinson and Ben Lippmeier. Machine fusion: merging merges, more or less. In
764 *Proceedings of the 19th International Symposium on Principles and Practice of Declarative*
765 *Programming*, pages 139–150, 2017.
- 766 **38** Amos Robinson and Alex Potanin. Pipit: Reactive systems in F*. In *Proceedings of the 8th*
767 *ACM SIGPLAN International Workshop on Type-Driven Development*, 2023.
- 768 **39** Indranil Saha and Suman Roy. A finite state analysis of time-triggered CAN (TTCAN)
769 protocol using Spin. In *2007 International Conference on Computing: Theory and Applications*
770 *(ICCTA’07)*, pages 77–81. IEEE, 2007.
- 771 **40** Ryan G Scott, Mike Dodds, Ivan Perez, Alwyn E Goodloe, and Robert Dockins. Trust-
772 worthy runtime verification via bisimulation (experience report). *Proceedings of the ACM on*
773 *Programming Languages*, 7(ICFP):305–321, 2023.
- 774 **41** Michael Short and Michael J Pont. Fault-tolerant time-triggered communication using CAN.
775 *IEEE transactions on Industrial Informatics*, 3(2):131–142, 2007.
- 776 **42** Joachim Zahnentferner, Dmytro Kaidalov, Jean-Frédéric Etienne, and Javier Díaz. Djed: a
777 formally verified crypto-backed autonomous stablecoin protocol. In *2023 IEEE International*
778 *Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–9. IEEE, 2023.