

Types for X10 Clocks

Francisco Martins
LaSIGE & University of Lisbon
Portugal
fmartins@di.fc.ul.pt

Vasco T. Vasconcelos
LaSIGE & University of Lisbon
Portugal
vv@di.fc.ul.pt

Tiago Cogumbreiro
LaSIGE & University of Lisbon
Portugal
cogumbreiro@di.fc.ul.pt

X10 is a modern language built from the ground up to handle future parallel systems, from multicore machines to cluster configurations. We take a closer look at a pair of synchronisation mechanisms: finish and clocks. The former waits for the termination of parallel computations, the latter allow multiple concurrent activities to wait for each other at certain points in time. In order to better understand these concepts we study a type system for a stripped down version of X10. The main result assures that well typed programs do not run into the errors identified in the X10 language reference, namely the `ClockUseException`. The study will open, we hope, doors to a more flexible utilisation of clocks in the X10 language.

1 Introduction

New high-level concurrency primitives are needed more than ever, now that multicore machines lay on our desks and laps. One such primitive is *clocks*, a generalisation of barriers introduced in the X10 programming language [2]. Barriers are a collective synchronisation mechanism common in Single Program Multiple Data (SPMD) programs [3, 12]. Distinct from synchronisation mechanisms like locks and monitors that let the programmer think about the access to a resource, barriers allow reasoning about process workflow. Clocks are a sophisticated form of barriers that feature dynamic sets of participants for more dynamic programming paradigms, and a two-phase synchronisation (or fuzzy barrier) for improved processor utilisation [5]. The construct is integrated in the X10 language with a promise that it cannot introduce deadlocks. Another primitive, finish, causes an activity to block (*i.e.*, to suspend its execution) until all its sub-activities have completed. Dynamic, unbounded spawning of activities and the finish construct enable fork/join parallelism. The fifth version of language Cilk introduced a specialised “work-stealing” scheduler algorithm that takes advantage of the fork/join model to yield highly efficient language [4]. This style of parallel programming and its run-time system (the work-stealing scheduler) were then incorporated in mainstream languages, *e.g.*, the fork/join framework proposed for Java 7 [6].

Even though the X10 language specification [8] provides a clear, plain English, description of the intended semantics (and properties) of the language, and a formalisation of the semantics [9] allows to prove a deadlock freedom theorem, we decided to investigate a simpler setting in which similar results could be obtained. The aim is not only to obtain a progress property for typable programs based on a simple type system, but also to hopefully provide for clock-safe extensions of the X10 language itself.

Towards this end, we have stripped X10 from most of its features, ending up with a simple concurrent language equipped with finish and the full functionality of X10 clocks, which we call “X10 restricted to clocks,” $X10|_{\text{clocks}}$ for short. For this language we have devised a simple operational semantics with thread (or activity as called in X10) local and global views of (heap allocated) clocks. We have also crafted a simple type system, based on singleton types, drawing expertise from previous work on low-level programming languages [13]. Typable programs are exempt from clock related errors; we conjecture that typable programs enjoy a form of progress property.

Saraswat and Jagadeesan study the X10 programming model, by presenting a formal model of the language that includes clocks, `async/finish`, conditional atomic blocks, and a hierarchic shared memory [9]. The authors formalise the semantics of the language with a small-step operational semantics, define a bisimulation, and establish that X10 programs without conditional atomic blocks do not deadlock. Lee and Palsberg present a core model for X10, an imperative language augmented with `async/finish` and atomic blocks suited for inter-procedural analysis through type inference [7]. The authors present an operational semantics based on [9] and a type system that identifies may-happen-parallelism, further explored in [1]. Java features, since version 5, cyclic barriers, `java.util.concurrent.CyclicBarrier`. Unlike X10 clocks these barriers have a fixed set of participants, defined at initialisation time.

This paper constitutes an archival version (post-proceedings) of a workshop paper produced in late 2009. Three of the language extensions proposed at that time have been incorporated in the X10 language definition [8], further discussed in Section 4. They are:

- Aliasing, introduced in version 2.01, January 2010. Prior versions imposed a restriction on how clock values could be aliased, “The initializer for a local variable declaration of type `Clock` must be a new clock expression. Thus X10 does not permit aliasing of clocks” (page 123, version 2.00).
- Clocks can be transmitted to a spawned activity if they are created in the scope of the enclosing `finish`, introduced in version 2.05, July 2010. Before we could read, “While executing `S` [the body of a `finish`], an activity must not spawn any `clocked` `asyncs`”¹ (page 141, version 2.04). Shirako *et al.* included this extension even prior to our work [10].
- Resume state inheritance, introduced in version 2.05, July 2010. Resumed clocks can be transmitted to forked activities. Prior to version 2.05 we could read, “It is a static error if any activity has a potentially live execution path from a `resume` statement on a clock `c` to a `async` `spawn` statement (which registers the new activity on `c`) unless the path goes through a `next` statement” (page 153, version 2.04).

In summary, the contributions of this work are:

- a simple operational semantics for activities, `finish`, and clocks that allows to better understand these constructs,
- a type systems allowing to prove safety and progress properties (alternative to the constraint-based system [9]), and
- the promise of a more flexible utilization of the clock constructs.

The rest of this paper presents the syntax in Section 2, the (operational) semantics and the notion of run-time errors in Section 3, the type system and some examples in Section 4, and the main result in Section 5. We conclude, in Section 6, discussing an alternative model for the semantics and pointing directions for future work.

2 Syntax

Object-oriented and type-safe, X10 provides for support for concurrency, parallelism, and distribution. Of particular interest to us is the `finish` synchronisation mechanism that waits for the termination of parallel computations and the clock primitive that allows forcing multiple concurrent activities to wait for each other at certain points in time.

¹In X10 a `clocked` `async` corresponds to an activity registered with at least one clock upon creation.

$e ::=$	<i>Expressions</i>	$v ::=$	<i>Values</i>
v	value	x	variable
async $\vec{v} e$	fork activity	$()$	unit
makeClock	new clock		
drop v	deregister from v		
finish e	wait to terminate		
next	advance phase		
resume v	ready to advance on v		
let $x = e$ in e	local declaration		

Figure 1: Top-level syntax of $X10|_{\text{clocks}}$

The top-level, or programmer's, language we address, $X10|_{\text{clocks}}$ (X10 restricted to clocks and finish), is a subset of the X10 language, generated by the grammar in Figure 1, and relies on a base set of *variables* ranged over by x . An $X10|_{\text{clocks}}$ program is an expression e that can operate on activities, clocks, or unit $()$ values. To construct programs we compose expressions through the standard let construct **let** $x = e_1$ **in** e_2 , which binds variable x to the result of expression e in the scope of expression e' .

Below we present an example program with the purpose of illustrating the syntax and informally presenting the semantics of the language. The example is composed of three activities an outermost activity a_1 , defined from line 2 to line 16, an inner activity a_2 , spawned at line 4, and another inner activity a_3 , spawned at line 5 and lasting until line 11. Along the example we make use of the derived expression $e_1; e_2$ that abbreviates **let** $x = e_1$ **in** e_2 , where x not free in e_2 .

```

// activity a1
finish
  let x = makeClock in (
    finish async e0; // activity a2
    async x ( // activity a3
      e1;
      resume x;
      e2;
      next;
      e3;
      drop x);
    e4;
    resume x;
    next;
    e5;
    drop x)

```

Activities can be *registered* with zero or more clocks and may share clocks with other activities. A clock can thus count with zero or more different registered activities, which are also called *participants*. When an activity a is registered with clock x , we say that x is a clock *held* by a . Activities may only register themselves with clocks via two different means: when they explicitly create a clock (line 3, **makeClock**, creates a clock and registers activity a_1 with the clock), and when they inherit clocks from

its parent activity in a spawning operation (line 5 spawns activity a_3 and registers it with the clock associated with variable x). Expression **drop** deregisters an activity from a clock (line 11, **drop** x deregisters activity a_3 from clock x , whereas line 16 deregisters activity a_1 from x). Activities are disallowed to manipulate clocks they are not registered with.

In $X10|_{\text{clocks}}$, an activity synchronises via two different methods: by waiting for every activity spawned by expression e to terminate (line 4 only terminates when activity a_2 and all its sub-activities terminate), and by waiting for its held clocks to *advance a phase* (activity a_3 waits at line 9, whereas activity a_1 waits at line 14). X10 distinguishes between *local* and *global* termination of an expression. Local termination of an expression corresponds to concluding its evaluation (reducing to a value v). An expression terminates globally when it terminates locally and each activity spawned by the expression has also terminated globally. Expression **finish** e converts the global termination of expression e into a local termination (line 2 waits for the expression in lines 3–16 to terminate, meaning that it waits as well for the termination of activities a_2 and a_3 ; line 4 waits for activity a_2 (the result of the evaluation of **async** e_0) to terminate before launching the sub-activity a_3 , in line 5). The second method to synchronise activities is using clocks. Groups of activities, defined by the participants of a clock, evaluate concurrently until they reach the end of a phase. When every participant of the group reaches the end of the phase, then all move to the next phase, while still executing concurrently. Phases are delimited by expression **next**; activities evaluate this expression to mark the end of a phase (activities a_1 and a_3 synchronise at lines 9 and 14).

An activity can inform all other participants of a clock x that it has completed its phase by using an expression of the form **resume** x , thus making clocks act as fuzzy barriers [5] (line 7). Expression **resume** can be viewed as an optimisation to diminish contention upon advancing a phase: it allows activities blocked on **next** to cease waiting for such activities (which can become at most one phase behind the clock's phase). In the example, expression e_2 might execute at the same time as expression e_5 , since activity a_3 may trigger (at line 5) activity a_1 to advance clock x (blocked at line 11), thus evaluating expressions e_2 and e_5 in parallel. If we omit expression **resume** x from this example, then expressions e_2 and e_5 cannot evaluate in parallel.

Clocks can be implemented via a (sophisticated) n -ary synchronisation mechanism that includes a natural number representing its *global phase*, initially set to zero. Advancing a clock's phase amounts to incrementing its global phase when every registered activity has *quiesced*; an activity is quiescent on a clock x after performing a **resume** x . An activity resumes all its held clocks together by evaluating **next** and suspends itself until these clocks become ready to advance to the next phase.

3 Operational Semantics

This section describes the operation semantics of our language, a possible execution of the example described in Section 2, and the notion of run-time errors.

Operational Semantics Figure 2 depicts the run-time syntax of our language. The run-time system relies on one additional set, *clock names* (or heap addresses, since clocks are the only data structures we allocate in the heap), ranged over by c . A state S of an $X10|_{\text{clocks}}$ computation comprises a shared heap H and a set of named activities A that run concurrently. Activity names, l , are taken from the set of variables introduced in Section 2. The heap stores clock values h , triples comprising a natural number p representing its global **phase**, a set R with the registered activities, and another set Q with the **quiesced** activities. These sets keep track of the activities that synchronise in the clock (R) and the activities that are

$S ::= H;A$	<i>States</i>	$e ::= \dots$	<i>Expressions</i>
$H ::= \{c_1 : h_1, \dots, c_n : h_n\}$	<i>Heaps</i>	$ \text{ join } l$	join activity
$A ::= \{l_1 : a_1, \dots, l_n : a_n\}$	<i>Sets of named activities</i>		
$h ::= \langle p, R, Q \rangle$	<i>Clock values</i>	$v ::= \dots$	<i>Values</i>
$R, Q ::= \{l_1, \dots, l_n\}$	<i>Sets of activity names</i>	$ c$	clock
$a ::= (V, e, A)$	<i>Activities</i>		
$V ::= \{c_1 : p_1, \dots, c_n : p_n\}$	<i>Clocks' local views</i>	$p ::= 0 \mid 1 \mid 2 \mid \dots$	<i>Phases</i>

Figure 2: Run-time syntax of $X10|_{\text{clocks}}$

ready to advance the clock to the next phase (Q). Set difference $R \setminus Q$ identifies the activities yet to make progress on a clock; the clock phase only advances when all activities have quiesced (when $R \setminus Q = \emptyset$, equivalently $R = Q$). The set of registered activities R also allows to enforce that an activity only operates on registered clocks.

An activity a is composed of a set of clocks' local views V , an expression e under execution, and a set of sub-activities A . Each activity has its own perception of the global phase of a clock; the clocks' local view V is a map from clock names to natural numbers describing the local phase. The global phase of a clock and that local to one of its activities may diverge in case the activity issues a **resume** on the clock. Only when the activity issues a **next**, the local view of the clock and the global phase become in sync. At anytime, a clock's local view is at most one phase behind the global phase. Notice that an activity is itself a tree of activities, since each activity holds a set of (named) sub-activities. When evaluating an expression **finish** e , the activity starts sub-activities for evaluating expression e and all activities spawned by e . Otherwise, activities have no sub-activities.

We augment the syntax of expressions at run-time with **join** l . Expression **join** l results from evaluating **finish** e ; label l identifies the activity that executes the body e of the **finish** expression and that produces the resulting value of the **finish** e expression.

We present the small step reduction rules for $X10|_{\text{clocks}}$ in Figures 3 and 4. Reduction for activities (Figure 3), $H; a \rightarrow_l H'; A; a'$, operates on a heap H and an activity a , and produces a possible different heap H' , a set A of activities spawned during the evaluation of the expression in a , and a new activity a' . Label l is the name of activity a .

Rule R-ASYNC is the only rule that affects the set of spawned activities A . The programmer specifies a list of clocks \vec{c} on which the new activity is to be registered with. The newly created activity (named l') is added to the set R of activities registered with clocks \vec{c} , and stored in the heap. The result of spawning an activity is the unit value $()$. The created activity is composed of a clock view holding, for each clock c in \vec{c} , a copy of the global phase p , an expression e to be evaluated, and an empty set of sub-activities. Syntax $H\{c : h\}$ describes a heap H' with a distinguished entry $c : h$; formally $H'(c) = h$ if $c = c'$ else $H(c)$. The new activity inherits each clock c quiescence property, *i.e.*, if l is quiescent on clock c so is l' ($Q' = \text{if } l \in Q \text{ then } Q \cup \{l'\} \text{ else } Q$).

The language specification reads “Clocks are created using a factory method on `x10.lang.Clock`. The current activity is automatically registered with the newly created clock” [8, page 207]. Expression **makeClock** creates a new clock in the heap with initial phase 0, with l as the only registered activity, and with no resumed activities, $\langle 0, \{l\}, \emptyset \rangle$. The activity creating the clock maintains a local clock

$$\begin{array}{c}
\frac{\{\vec{c}\} \subseteq \text{dom}V \quad l' \text{ is fresh} \quad Q' \triangleq \text{if } l \in Q \text{ then } Q \cup \{l'\} \text{ else } Q}{H\{c: \langle p, R, Q \rangle\}_{c \text{ in } \vec{c}}; (V, \mathbf{async} \vec{c} e, A) \rightarrow_l H\{c: \langle p, R \cup \{l'\}, Q' \rangle\}_{c \text{ in } \vec{c}}; \{l': (\{c: p\}_{c \text{ in } \vec{c}}, e, \emptyset)\}; (V, (), A)} \quad (\text{R-ASYNC}) \\
\\
\frac{c \text{ is fresh}}{H; (V, \mathbf{makeClock} A') \rightarrow_l H\{c: \langle 0, \{l\}, \emptyset \rangle; \emptyset; (V\{c: 0\}, c, A')} \quad (\text{R-MAKE}) \\
\\
\frac{Q' \triangleq \text{if } p = V(c) \text{ then } Q \cup \{l\} \text{ else } Q \quad l \notin Q}{H\{c: \langle p, R, Q \rangle\}; (V, \mathbf{resume} c, A) \rightarrow_l H\{c: \langle p, R, Q' \rangle\}; \emptyset; (V, (), A)} \quad (\text{R-RESUME}) \\
\\
\frac{C_1 \triangleq \{c \mid V(c) = p, H(c) = \langle p, R, R \rangle\} \quad C_2 \triangleq \{c \mid V(c) = p, H(c) = \langle p+1, -, - \rangle\} \quad C_1 \cup C_2 = \text{dom}V}{H; (V, \mathbf{next}, A) \rightarrow_l H\{c: \langle p+1, R, \emptyset \rangle\}_{c \in C_1}; \emptyset; (\{c: V(c) + 1\}_{c \in V}, (), A)} \quad (\text{R-NEXT}) \\
\\
\frac{c \in \text{dom}V \quad H' \triangleq \text{if } R = \{l\} \text{ then } H \setminus \{c\} \text{ else } H\{c: \langle p, R \setminus \{l\}, Q \setminus \{l\} \rangle\}}{H\{c: \langle p, R, Q \rangle\}; (V, \mathbf{drop} c, A) \rightarrow_l H'; \emptyset; (V \setminus \{c\}, (), A)} \quad (\text{R-DROP}) \\
\\
\frac{l_0 \text{ is fresh}}{H; (V, \mathbf{finish} e, A) \rightarrow_l H'; \emptyset; (V, \mathbf{join} l_0, A\{l_0: (V, e, \emptyset)\})} \quad (\text{R-FINISH}) \\
\\
H; (V, \mathbf{join} l_0, \{l_0: (\emptyset, v_0, \emptyset), \dots, l_n: (\emptyset, v_n, \emptyset)\}) \rightarrow_l H'; \emptyset; (\emptyset, v_0, \emptyset) \quad (\text{R-JOIN})
\end{array}$$

Figure 3: Reduction rules for activities $H; a \rightarrow_l H; A; a$

$$\begin{array}{c}
H; A\{l: (V, \mathbf{let} x = v \text{ in } e, A')\} \rightarrow H; A\{l: (V, e[v/x], A')\} \quad (\text{R-LET-VAL}) \\
\\
\frac{l \in \text{dom}H \quad H; (V, e, A) \rightarrow_l H'; A''; (V', e', A')}{H; A''\{l: (V, \mathbf{let} x = e \text{ in } e'', A)\} \rightarrow H'; A''\{l: (V', \mathbf{let} x = e' \text{ in } e'', A')\}, A'''} \quad (\text{R-LET}) \\
\\
\frac{H; A' \rightarrow H'; A''}{H; A\{l: (V, e, A')\} \rightarrow H'; A\{l: (V, e, A'')\}} \quad (\text{R-ACTIVITY})
\end{array}$$

Figure 4: Reduction rules for states $H; A \rightarrow H; A$

view $\{c: 0\}$ stored in V . Regarding expression **resume** the language specification reads “An activity may wish to indicate that it has completed whatever work it wishes to perform in the current phase of a clock c it is registered with, without suspending altogether” (page 208). Rule R-RESUME asserts that when the l -labelled activity issues a **resume** c , its label is recorded in the set of resumed activities R if the clock local phase is in sync with the clock global phase ($p = V(c)$); otherwise, the effect of the expression is discarded ($p \neq V(c)$), since the clock has already advanced to the next phase. An activity may only perform a **resume** operation per clock per phase ($l \notin Q$).

The language specification reads “Execution of this statement [**next**] blocks until all the clocks that the activity is registered with (if any) have advanced. (The activity implicitly issues a resume on all clocks it is registered with before suspending.) [. . .]. An activity blocked on next resumes execution once it is marked for progress by all the clocks it is registered with” (page 209). In our model, for simplicity’s sake, the programmer must issue a **resume** on all held clocks before expression **next**. As rule R-NEXT states, expression **next** blocks the activity until all clocks have been resumed (C_1) or have

already advance their phases (C_2). Notice that when activities are waiting on a clock c , the clock can be in one of three states: (a) there are non-quiescent activities on the clock and c is neither a member of C_1 nor of C_2 ; (b) all registered activities are quiescent on the clock, and so c is a member of C_1 ; (c) the clock has advanced to the next phase thus becoming a member of C_2 . When an activity advances a clock global phase, it stops being a member of set C_1 and becomes a member of set C_2 for the remaining activities waiting on that clock. Since rule R-NEXT only updates the clock phase of those clocks belonging to C_1 ($H\{c: \langle p+1, R, \emptyset \rangle\}_{c \in C_1}$) it ensures that the global clock state is updated only once.

The language specification reads “An activity may drop a clock by executing $c.\text{drop}()$. The activity is no longer considered registered with this clock” (page 209). With expression **drop** c , the l -labelled activity cedes its control over clock c : we remove c from clock view V , and remove activity identifier l from both sets R and Q . Two consequences of dropping a clock c are: a) activities waiting on clock c are no longer blocked because of this activity; b) when executing a **next** expression, this activity no longer waits for clock c . In case l is the only activity registered with clock c , it is safe to deallocate the clock, so that the clock’s heap space can be reclaimed without resorting to garbage collection. The language specification reads “An activity A executes finish S by executing S and then waiting for all activities spawned by S (directly or indirectly [. . .]) to terminate” (page 196). Expression **finish** e creates a child activity and evaluates into expression **join** l_0 (rule R-FINISH), which in turn blocks while there exist sub-activities running. When all sub-activities have reduced to a value, activity l (**join** l_0) evaluates to the value in its sub-activity l_0 and garbage collects all other sub-activities (rule R-JOIN).

The reduction for states (Figure 4), $S \rightarrow S'$, allows for non-deterministic choice of which activity l to evaluate (rule R-ACTIVITY), capturing the concurrency present in X10 computations. We evaluate the let binding from left-to-right (rule R-LET), when the left-hand-side expression becomes a value, we substitute this value for variable x in the continuation expression e (rule R-LET-VAL).

Example Recall the example from Section 2. Consider a loading function that sets up the initial state from a given expression, which in this case is S_0 , an empty heap and an activity evaluating the code in the example under a dummy **let**. State S_0 reduces in two steps, using rules R-FINISH and R-LET. The grey boxes highlight a redex and also the corresponding contractum.

$$S_0 = \emptyset; \{l_1 : (\emptyset, \text{let } z = \underbrace{\text{finish let } x = \text{makeClock in } (\text{finish } (\text{async } e_0); e_6)}_{\text{Example from Section 2}} \text{ in } (), \emptyset)\}$$

Example from Section 2

where e_6 is $(\text{async } x (e_1; \text{resume } x; e_2; \text{next}; e_3; \text{drop } x)); \text{resume } x; e_4; \text{next}; e_5; \text{drop } x$. We perform two reduction steps to illustrate the effect of expression **finish** on the sub-activities of l_1 .

$$\emptyset; \{l_1 : (\emptyset, \text{let } z = \text{join } l_2 \text{ in } (), \{l_2 : (\emptyset, \text{let } x = \text{makeClock in } (\text{finish } (\text{async } e_0); e_6), \emptyset)\})\}$$

From this state on, while **join** remains blocked, we apply rule R-ACTIVITY to evaluate the child activities of l_1 . We perform four further reduction steps (R-ACTIVITY, R-LET, R-MAKE, and R-LET-VAL) and observe how expression **makeClock** updates the heap and the clock view of activity l_2 .

$$\{c : \langle 0, \{l_2\}, \emptyset \rangle\}; \{l_1 : (\emptyset, \text{let } z = \text{join } l_2 \text{ in } (), \{l_2 : (\{c : 0\}, \text{let } x = c \text{ in } (\text{finish } (\text{async } e_0); e_6), \emptyset)\})\}$$

The non-determinism of our semantics now allows for various different reductions. A possible outcome of a (multi-step) reduction is

$$\{c : \langle 0, \{l_2, l_3\}, \emptyset \rangle\}; \{l_1 : (\emptyset, \text{let } z = \text{join } l_2 \text{ in } (), \{l_2 : (\{c : 0\}, (\text{resume } c; \text{next}; e_5; \text{drop } c), \emptyset), \\ l_3 : (\{c : 0\}, (\text{resume } c; e_2; \text{next}; e_3; \text{drop } c), \emptyset)\})\}$$

$$\begin{array}{l}
H;A\{l: (V, \mathbf{let} \ x = \mathbf{async} \ \vec{c} \ e \ \mathbf{in} \ e', A')\} \in \text{Error} \quad \text{if } \vec{c} \not\subseteq \text{dom}H \text{ or } c \notin \text{dom}V \quad (\text{E-ASYNC}) \\
H\{c: \langle _, _, Q \rangle\}; A\{l: (V, \mathbf{let} \ x = \mathbf{resume} \ c \ \mathbf{in} \ e, A')\} \in \text{Error} \quad \text{if } l \in Q \text{ or } c \notin \text{dom}H \text{ or } c \notin \text{dom}V \\
\hspace{20em} (\text{E-RESUME}) \\
H;A\{l: (V, \mathbf{let} \ x = \mathbf{drop} \ c \ \mathbf{in} \ e, A')\} \in \text{Error} \quad \text{if } c \notin \text{dom}H \text{ or } c \notin \text{dom}V \quad (\text{E-DROP}) \\
H;A\{l: (V, \mathbf{let} \ x = \mathbf{next} \ \mathbf{in} \ e, A')\} \in \text{Error} \quad \text{if } V(c) = p, H(c) = (p, _, Q), \text{ and} \\
\hspace{15em} l \notin Q, \text{ for some } c \quad (\text{E-NEXT1}) \\
H;A\{l: (V, \mathbf{let} \ x = \mathbf{next} \ \mathbf{in} \ e, A')\} \in \text{Error} \quad \text{if } c \in \text{dom}V \text{ and } c \notin \text{dom}H, \text{ for some } c \\
\hspace{20em} (\text{E-NEXT2}) \\
H;A\{l: (V, v, _)\} \in \text{Error} \quad \text{if } V \neq \emptyset \quad (\text{E-ACT}) \\
\frac{H;A' \in \text{Error}}{H;A\{l: (_, _, A')\} \in \text{Error}} \quad (\text{E-ACT-SET})
\end{array}$$

Figure 5: The set Error of run-time errors

We now illustrate the case when activities l_2 and l_3 are evaluating expressions e_2 and e_5 concurrently.

$$\begin{array}{l}
\text{R-ACTIVITY, R-LET} \rightarrow \text{R-RESUME} \xrightarrow{l_2} \text{R-ACTIVITY, R-LET-VAL} \rightarrow \\
\{c: \langle 0, \{l_2, l_3\}, \{l_2\} \rangle\}; \{l_1: (\emptyset, \mathbf{let} \ z = \mathbf{join} \ l_2 \ \mathbf{in} \ (), \{l_2: (\{c: 0\}, (\mathbf{next}; e_5; \mathbf{drop} \ c), \emptyset), \\
\hspace{15em} l_3: (\{c: 0\}, (\mathbf{resume} \ c; e_2; \mathbf{next}; e_3; \mathbf{drop} \ c), \emptyset))\} \\
\text{R-ACTIVITY, R-LET} \rightarrow \text{R-RESUME} \xrightarrow{l_3} \\
\{c: \langle 0, \{l_2, l_3\}, \{l_2, l_3\} \rangle\}; \{l_1: (\emptyset, \mathbf{join} \ l_2, \{l_2: (\{c: 0\}, (\mathbf{next}; e_5; \mathbf{drop} \ c), \emptyset), \\
\hspace{15em} l_3: (\{c: 0\}, (\emptyset; e_2; \mathbf{next}; e_3; \mathbf{drop} \ c), \emptyset))\} \\
\text{R-ACTIVITY, R-LET} \rightarrow \text{R-NEXT} \xrightarrow{l_2} \\
\{c: \langle 1, \{l_2, l_3\}, \emptyset \rangle\}; \{l_1: (\emptyset, \mathbf{let} \ z = \mathbf{join} \ l_2 \ \mathbf{in} \ (), \{l_2: (\{c: 1\}, (\emptyset; e_5; \mathbf{drop} \ c), \emptyset), \\
\hspace{15em} l_3: (\{c: 0\}, (\emptyset; e_2; \mathbf{next}; e_3; \mathbf{drop} \ c), \emptyset))\}
\end{array}$$

After activity l_3 evaluates **resume** c , activity l_2 , which is blocked evaluating **next**, progresses, thus allowing expressions e_2 and e_5 to execute in parallel. Notice that activity l_2 remains in phase 0, while activity l_3 is in phase 1.

Run-time errors Run-time errors is the smallest set Error of states generated by the rules in Figure 5. The notion is consistent with all the conditions documented to raise exception `ClockUseException`, as discussed in the X10 language specification report [8]. The type system we present in Section 4 allow us to reject, at compile time, programs that can potentially throw a `ClockUseException`.

During an **async** operation, an activity cannot transmit unregistered clocks through its first argument (rule E-ASYNC). Similarly, activities can only perform **resume** or **drop** operations on clocks they are registered with (rules E-RESUME and E-DROP). In particular, it constitutes an error for an activity to drop a clock twice, or to resume a clock more than once (for the same phase) or after dropping it. We achieved a fine grained control over the clocks an activity is registered with. Specifically, it is possible to devise, at compile time, whether an activity resumed or dropped all of its held clocks, as manifest

$$\tau ::= \mathbf{unit} \mid \mathbf{clock}(\alpha) \qquad \text{Types}$$

Figure 6: Syntax of types

$$\begin{array}{c}
\mathcal{R}, \alpha \vdash \mathbf{clock}(\alpha) \quad \mathcal{R} \vdash \mathbf{unit} \qquad \text{(T-WF-C, T-WF-U)} \\
\frac{\mathcal{R} \vdash \tau}{\Gamma, x: \tau; \mathcal{R} \vdash x: \tau} \quad \Gamma, c: \mathbf{clock}(\alpha); \mathcal{R}, \alpha \vdash c: \mathbf{clock}(\alpha) \quad \Gamma; \mathcal{R} \vdash (): \mathbf{unit} \\
\text{(T-VAR, T-CLOCK-REF, T-UNIT)} \\
\frac{\Gamma; \mathcal{R} \vdash v_1: \mathbf{clock}(\alpha_1) \quad \cdots \quad \Gamma; \mathcal{R} \vdash v_n: \mathbf{clock}(\alpha_n) \quad \alpha_i \neq \alpha_j, \text{ if } i \neq j \quad \alpha_i \text{ not in } \Gamma}{\Gamma; \mathcal{R} \vdash v_1 \dots v_n: \{\alpha_1, \dots, \alpha_n\}} \\
\text{(T-CLOCK-SEQ)}
\end{array}$$

Figure 7: Typing rules for values and for well-formed types

from our typing rules later. Therefore, it constitutes an error when an activity evaluates a **next** expression before resuming all its clocks (rule E-NEXT1). It is also an error when upon evaluating a **next** there is a held clock that is not in the heap (rule E-NEXT2). Furthermore, an activity cannot evaluate to a value without dropping all of its clocks (rule E-ACT). Rule E-ACT-SET allows error propagation from sub-activities.

Earlier versions of the language specification report (until version 2.05) included two additional error conditions we quote:

- “It is a static error if any activity has a potentially live execution path from a **resume** statement on a clock c to a **async spawn** statement (which registers the new activity on c) unless the path goes through a **next** statement” (page 153, version 2.04). (See example 2, Section 4.);
- “While executing S [the body of a **finish**], an activity must not spawn any **clocked** **asyns**. (Asyns spawned during the execution of S may spawn **clocked** **asyns**.)” (page 141, version 2.04).

Our type system guarantees soundness in presence of these two conditions.

4 Type System

This section presents a type system that uses *singleton types* to track clock usage throughout a program.

For types we rely on an additional base set of singleton types ranged over by α . The syntax of types, depicted in Figure 6, introduces the type **unit** of unit values, and the type **clock**(α) of a *particular* clock. We assign a different type to each clock in order to ensure the correct usage of the clock constructs within a program.

The type system for $X10|_{\text{clocks}}$ programs is defined in Figures 7 and 8. A typing Γ is a map from variables (or activity labels) and clocks to types. We write $\text{dom } \Gamma$ for the domain of Γ . When $x \notin \text{dom } \Gamma$ we write $\Gamma, x: \tau$ for the typing Γ' such that $\text{dom } \Gamma' = \text{dom } \Gamma \cup \{x\}$, $\Gamma'(x) = \tau$, and $\Gamma'(y) = \Gamma(y)$ for $y \neq x$. The type system also uses sets of singleton types, ranged over by \mathcal{R} , for **registered** clocks, and \mathcal{Q} , for **quiescent** clocks.

$$\begin{array}{c}
\frac{\Gamma; \mathcal{R} \vdash v: \tau}{\Gamma; \mathcal{R}; \mathcal{Q} \vdash v: (\tau, \mathcal{R}, \mathcal{Q})} \quad \frac{\alpha \text{ is fresh}}{\Gamma; \mathcal{R}; \mathcal{Q} \vdash \mathbf{makeClock}: (\mathbf{clock}(\alpha), \mathcal{R} \cup \{\alpha\}, \mathcal{Q})} \quad \text{(T-VALUE, T-MAKE)} \\
\frac{\Gamma; \mathcal{R} \vdash v: \mathbf{clock}(\alpha) \quad \alpha \notin \mathcal{Q}}{\Gamma; \mathcal{R}; \mathcal{Q} \vdash \mathbf{resume } v: (\mathbf{unit}, \mathcal{R}; \mathcal{Q} \cup \{\alpha\})} \quad \frac{\Gamma; \mathcal{R} \vdash v: \mathbf{clock}(\alpha)}{\Gamma; \mathcal{R}; \mathcal{Q} \vdash \mathbf{drop } v: (\mathbf{unit}, \mathcal{R} \setminus \{\alpha\}, \mathcal{Q} \setminus \{\alpha\})} \\
\text{(T-RESUME, T-DROP)} \\
\frac{\Gamma; \mathcal{R} \vdash \vec{v}: \mathcal{R}' \quad \Gamma; \mathcal{R}'; \mathcal{Q} \cap \mathcal{R}' \vdash e: (-, \emptyset, \emptyset)}{\Gamma; \mathcal{R}; \mathcal{Q} \vdash \mathbf{async } \vec{v} e: (\mathbf{unit}, \mathcal{R}; \mathcal{Q})} \quad \frac{}{\Gamma; \mathcal{R}; \mathcal{R} \vdash \mathbf{next}: (\mathbf{unit}, \mathcal{R}; \emptyset)} \quad \text{(T-ASYNC, T-NEXT)} \\
\frac{\Gamma; \emptyset; \emptyset \vdash e: (\tau, \emptyset, \emptyset)}{\Gamma; \mathcal{R}; \mathcal{Q} \vdash \mathbf{finish } e: (\tau, \mathcal{R}; \mathcal{Q})} \quad \text{(T-FINISH)} \\
\frac{\Gamma; \mathcal{R}; \mathcal{Q} \vdash e_1: (\tau, \mathcal{R}', \mathcal{Q}') \quad \Gamma, x: \tau; \mathcal{R}'; \mathcal{Q}' \vdash e_2: (\tau', \mathcal{R}'', \mathcal{Q}'')}{\Gamma; \mathcal{R}; \mathcal{Q} \vdash \mathbf{let } x = e_1 \mathbf{ in } e_2: (\tau', \mathcal{R}'', \mathcal{Q}'')} \quad \text{(T-LET)}
\end{array}$$

Figure 8: Typing rules for expressions

The typing rules for values and for well formed types (Figure 7) are simple to follow. Well-formedness for clock types (rule T-WF-C) ensures that activities only make use of clocks they are registered with. Rule T-CLOCK-SEQ ensures that different clocks (as those in the heap) have distinct singleton clock types, a property that is crucial for establishing type safety. For typing expressions we use a type system (Figure 8) that records the changes made to the set of registered clocks, either by creating or dropping clocks, and to the set of quiescent clocks (using **resume** and **next**) of an expression. Typing judgements are of the form $\Gamma; \mathcal{R}; \mathcal{Q} \vdash e: (\tau, \mathcal{R}', \mathcal{Q}')$ meaning that expression e is well typed assuming the types for the free identifiers in Γ , the registered clocks in \mathcal{R} , and the quiescent clocks in \mathcal{Q} . The type of an expression is a triple recording the type τ of its value, as well as the registered \mathcal{R}' and the quiescent \mathcal{Q}' sets after execution of the expression.

Most typing rules are straightforward. When creating a clock (rule T-MAKE) we associate a new singleton type α with the clock and include it in set of clocks registered by the activity ($\mathcal{R} \cup \{\alpha\}$). Rule T-RESUME, which asserts that “an activity may invoke **resume**() only on a clock it is registered with, and has not yet dropped” (page 208, *vide* rule T-VAR), marks clock α as quiescent. Notice that a clock cannot be resumed more than once for the same phase ($\alpha \notin \mathcal{R}$), contrary to the language reference that reads, “Nothing happens if the activity has already invoked a resume on this clock in the current phase” (page 208). A **drop** v expression removes clock v from both the sets \mathcal{R} and \mathcal{Q} , thus the clock cannot be passed to new activities, be the target of a **resume** expression, or be dropped again.

For expression **async** $\vec{v} e$, the language reference reads “Starts a new activity, initially registered with clocks $[\vec{v}]$, and running $[e]$. The activity running this code must be registered on those clocks” (page 207, w.r.t. rule T-ASYNC). Rule T-ASYNC asserts that when an activity spawns another activity registered on a sequence of clocks, the quiescent property of the clocks is preserved by propagating the information about the quiescent clock α ($\mathcal{Q} \cap \mathcal{R}$). Moreover, the new activity must have dropped all its clocks upon termination, contrary to the language reference that reads, “All activities are automatically deregistered from all clocks they are registered with on termination (normal or abrupt)” (page 207). An activity cannot share a clock it does not hold, as noted in the language reference, “lacking that registration, cannot register a sub-activity on it [a clock] with **async**” (page 208). Expression **next** marks the end of a phase; it checks that all clocks have been resumed and clears the quiescent clocks for the new phase (rules T-NEXT).

The **finish** construct may interfere with clocks and cause programs to deadlock. In order to avoid such situations we prevent the body of a **finish** e expression (e) from accessing any clock already defined, thus eliminating (nested) dependencies between clocks and **finish**. Rule T-FINISH also forces e to unregister from all clocks it has created, and therefore **finish** e has no effect on registered and quiescent clocks. This follows the semantics of the current version of X10 that reads, “Inside of **finish**{S}, all **clocked** **asyncs** must be in the scope an **unclocked** **async**” (page 210). Refer to the examples below for further discussion on the deadlock problem. When typing a **let** expression (rule T-LET), its continuation e_2 is typed taking into consideration the effects produced by expression e_1 . The type of the **let** is that of e_2 , as usual.

We have deliberately deviated from the standard X10 semantics in three cases: **next**, **drop**, and **resume**. The reasons for such deviation are: (a) to illustrate the power of singleton types in keeping track of clocks, (b) to simplify the (operational and static) semantics, (c) to enforce a programming discipline that may avoid potential bugs, and (d) because the compiler has enough information to suggest, or automatically insert, code fixes (*e.g.*, by enumerating the clocks that need to be dropped before a **next**; see examples below).

Below we discuss a few $X10|_{\text{clocks}}$ programs and the semantic guarantees our type system enforces. We decorate the examples with the typing assumptions $(\Gamma, \mathcal{R}, \mathcal{Q})$ holding for each expression.

Example 1: Aliasing Our first example concerns clock aliasing, only introduced in X10 in version 2.01. The example may read a bit trivial but illustrates more sophisticated aliasing situations, derived for example from procedure calls. Clearly a type system with linear control like the one we are going to present allows to relieve such a restriction.

```

// activity a1
let x = makeClock in ( // {x: clock(α)}, {α}, ∅ 1
  async x ( // activity a2 // {x: lock(α)}, {α}, ∅ 2
    let y = x in ( // {x: clock(α), y: clock(α)}, {α}, ∅ 3
      resume x; // {x: clock(α), y: clock(α)}, {α}, {α} 4
      drop y; // {x: clock(α), y: clock(α)}, ∅, ∅ 5
    ); // {x: clock(α)}, {α}, ∅ 6
  drop x) // ∅, ∅, ∅ 7

```

In our case the code is typable, assigning the same singleton type **clock**(α) to both x and y . Upon introducing variable x (line 2), it gets assigned type **clock**(α) (*vide* rules T-LET and T-MAKE). Variable y (line 4) gets assigned type **clock**(α), the type of x (again using rules T-LET and T-MAKE).

Example 2: Resume state inheritance Our second example deals with a restriction the language reference disallowed up until version 2.04, “A potentially live execution path from a **resume** statement on a clock c to an **async** spawn statement” (page 153, version 2.04). Our operational semantics allows the forked activity to inherit the *resume state* (resumed/not resumed) of the parent activity (*vide* rule R-ASYNC) and therefore preserves the quiescence property of clocks and avoids a race condition on the clock. Notice that the forked activity a_2 inherits the quiescent state of clock α (line 4), resumed at line 3.

```

// activity a1
let x = makeClock in ( // {x: clock(α)}, {α}, ∅ 1
  resume x; // {x: clock(α)}, {α}, {α} 2

```

```

async x ( // activity a2           // {x: clock( $\alpha$ ), { $\alpha$ }, { $\alpha$ }}           4
  next;           // {x: clock( $\alpha$ ), { $\alpha$ },  $\emptyset$ }           5
  drop x         // {x: clock( $\alpha$ ),  $\emptyset$ ,  $\emptyset$ }           6
);           // {x: clock( $\alpha$ ), { $\alpha$ }, { $\alpha$ }}           7
drop x)       // {x: clock( $\alpha$ ),  $\emptyset$ ,  $\emptyset$ }           8

```

Example 3: Race condition generated by not inheriting the resume state Describes a race condition triggered by clock synchronisation. Activity a_1 creates a clock x , starts a second activity a_2 registered with clock x that, in turn, resumes on x and starts a third activity a_3 also registered with x .

```

// activity a1           1
let x = makeClock in ( // {x: clock( $\alpha$ ), { $\alpha$ },  $\emptyset$ }           2
  async x ( // activity a2 // {x: clock( $\alpha$ ), { $\alpha$ },  $\emptyset$ }           3
    resume x; // {x: clock( $\alpha$ ), { $\alpha$ }, { $\alpha$ }}           4
    async x ( // activity a3 // {x: clock( $\alpha$ ), { $\alpha$ }, { $\alpha$ }}           5
      next; // {x: clock( $\alpha$ ), { $\alpha$ },  $\emptyset$ }           6
      drop x // {x: clock( $\alpha$ ),  $\emptyset$ ,  $\emptyset$ }           7
    ); // {x: clock( $\alpha$ ), { $\alpha$ }, { $\alpha$ }}           8
    next; // {x: clock( $\alpha$ ), { $\alpha$ },  $\emptyset$ }           9
    drop x // {x: clock( $\alpha$ ),  $\emptyset$ ,  $\emptyset$ }           10
  ); // {x: clock( $\alpha$ ), { $\alpha$ },  $\emptyset$ }           11
  resume x; // {x: clock( $\alpha$ ), { $\alpha$ }, { $\alpha$ }}           12
  next; // {x: clock( $\alpha$ ), { $\alpha$ },  $\emptyset$ }           13
  drop x) // {x: clock( $\alpha$ ),  $\emptyset$ ,  $\emptyset$ }           14

```

The race condition might occur because after a_2 resumes on x (line 4), either activity a_1 may advance clock x phase by executing **next** (line 13) or a_2 may register a new activity a_3 with x (line 5), blocking a_1 until activity a_3 executes its **next** instruction (line 6). By inheriting the resume status of clock x , activity a_3 does not block activity a_1 and the race condition disappears (*vide* rule R-ASYNC in Figure 3 and rule T-ASYNC in Figure 8).

Example 4: Resume after resume The next example deals with resuming after resuming, a pattern accepted in X10. Rule T-RESUME rejects the program below, since it is able to determine that clock x is resumed twice.

```

// activity a1           1
let x = makeClock in ( // {x: clock( $\alpha$ ), { $\alpha$ },  $\emptyset$ }           2
  resume x; // {x: clock( $\alpha$ ), { $\alpha$ }, { $\alpha$ }}           3
  resume x) // error: clock x already quiescent for activity a1 4

```

Example 5: Explicit drops Unlike X10, we have decided to explicitly deregister activities from clocks upon activity termination. Our type system keeps track of the clocks an activity is registered with, and rejects programs with activities that finish before deregistering from all its clocks. Clocks without registered activities can be safely garbage collected (*vide* rule R-DROP). The following example fails to type check, since the launched activity does not drop clock x . The compiler may easily suggest an appropriate fix: adding a **drop** x after the **next** instruction on line 5, or even automatically introduce such an instruction.

```

// activity a1
let x = makeClock in (
  async x ( // activity a2
    resume x;
    next
  );
drop x)

```

1
2
3
4
5
6
7

Example 6: Finish/async deadlock Finally, we discuss the interplay among **finish**, **async**, and clocks, which may cause programs to deadlock. The following program deadlocks because activity a_2 is waiting on **next** (line 6) for activity a_1 to advance on x , which is planned to occur at line 10, but a_1 is waiting on **finish** (line 3) for activity a_2 to terminate, so a_1 never reaches line 10 and the program deadlocks.

```

// activity a1
let x = makeClock in (
  finish
  async x ( // activity a2 // error: clock x is not in scope of finish
    resume x;
    next;
    drop x
  );
resume x;
next;
drop x)

```

1
2
3
4
5
6
7
8
9
10
11

The cause for deadlock is that activity a_2 is registered with a clock that is defined outside the enclosing **finish**: clock x is defined in line 2, whereas the **finish** expression extends from line 3 to line 8. Our type system rejects this program, because when typing a **finish** e expression we type check e in an environment with no registered clocks (*vide* rule R-FINISH).

Example 7: Clocked finish/clocked async Version 2.10 of the X10 language introduces keyword **clocked** to prefix **async** and **finish**.

In the most common case of a single clock coordinating a few behaviors, X10 allows coding with an implicit clock. [...] A **clocked finish** introduces a new clock. It executes its body in the usual way that a **finish** does—except that, when its body completes, the activity executing the clocked **finish** drops the clock, while it waits for asynchronous spawned **asyncs** to terminate. A **clocked async** registers its **async** with the implicit clock of the surrounding clocked **finish**. [...] Clocked finishes may be nested. The inner **clocked finish** operates in a single phase of the outer one.

This feature introduced in the language is purely “syntactic sugar,” which makes the common practice of creating a single clock and sharing it among activities simpler. The following two code listings present a program with and without the syntactic extension side-by-side. On the left column we show an activity written on a language that resembles X10 [8]. Remember that, in X10, expression **next** implicitly issues a **resume** and also that activities implicitly drop all clocks on exit. On the right column we find an equivalent activity written in our language. Activity a_1 spawns three activities a_2 , a_3 , and a_4 . Activities a_2 and a_3 share the same clock, say c_1 , whereas activity a_4 holds a different clock, say c_2 , but it

does not hold c_1 . In the first phase of clock c_1 the system executes concurrently e_1 , e_3 , and activity a_4 . In the second phase of clock c_1 , activity a_4 has terminated, and expressions e_2 and e_4 execute concurrently.

<pre> // activity a1 clocked finish (clocked async (// activity a2 e1 next; e2); clocked async (// activity a3 e3 next; e4) clocked finish (clocked async (// activity a4 e5 next; e4))) </pre>	<pre> // activity a1 finish //0,0,0 let c1 = makeClock in (// {c1: clock(α1)}, {α1}, 0 async c1 (// activity a2 // {c1: clock(α1)}, {α1}, 0 e1 resume c1; next; // {c1: clock(α1)}, {α1}, 0 e2 drop c1; // {c1: clock(α1)}, 0, 0); // {c1: clock(α1)}, {α1}, 0 async c1 (// activity a3 // {c1: clock(α1)}, {α1}, 0 e3 resume c1; next; // {c1: clock(α1)}, {α1}, 0 e4 drop c1; // {c1: clock(α1)}, 0, 0); // {c1: clock(α1)}, {α1}, 0 finish // {c1: clock(α1)}, 0, 0 let c2 = makeClock in (// {c1: clock(α1), // c2: clock(α2)}, {α2}, 0 async c2 (// activity a4 // {c1: clock(α1), // c2: clock(α2)}, {α2}, 0 e5 // {c1: clock(α1), resume c2; next; // c2: clock(α2)}, {α2}, 0 e4 // c2: clock(α2)}, {α2}, 0 drop c2 // {c1: clock(α1), c2: clock(α2)}, 0, 0); // {c1: clock(α1), c2: clock(α2)}, {α2}, 0 drop c2 // {c1: clock(α1), c2: clock(α2)}, 0, 0); // {c1: clock(α1)}, {α1}, 0 drop c1 // {c1: clock(α1)}, 0, 0) //0,0,0 </pre>
---	--

The activity (on the right) obtained by expanding **clocked finish** and **clocked async** is still typable with the type system we propose.

5 Main results

This section is dedicated to the study of the main result of our system, namely typing preservation and type safety for typable programs.

We are only interested in well-formed states (the rules in Figure 9 check whether a state is well-formed). A state is well formed if for each clock, the set of registered activities with the clock contains exactly those activities that can manipulate it. State

$$\{c: \langle -, \emptyset, - \rangle\}; \{l: (\{c: -\}, \text{let } - = \text{next in } -, -)\}$$

is ill formed, since activity l uses clock c and is not registered with c . The activity is able to advance c 's

$$\begin{array}{c}
\frac{c \vdash A: S_1 \quad c \vdash A': S_2}{c \vdash A, \{l: (V, _, A')\}: S_1 \cup S_2 \cup \text{dom } V \cap \{c\}} \quad c \vdash \emptyset: \emptyset \\
\text{(WF-ACT-CLOCK, WF-ACT-CLOCK-E)} \\
\frac{Q \subseteq S \subseteq \text{dom } \Gamma \quad c \vdash A: S \quad \Gamma; A \vdash H: \diamond}{\Gamma; A \vdash H, \{c: \langle _, S, Q \rangle\}: \diamond} \quad \Gamma; A \vdash \emptyset: \diamond \quad \text{(WF-HEAP, WF-HEAP-E)} \\
\frac{H(c_i) = \langle _, \mathcal{R}_i, _ \rangle \quad l \in \mathcal{R}_i \quad H \vdash A: \diamond \quad H \vdash A': \diamond}{H \vdash A, \{l: (\{c_1: _, \dots, c_n: _ \}, _, A')\}: \diamond} \quad H \vdash \emptyset: \diamond \\
\text{(WF-ACT-SET, WF-ACT-SET-E)} \\
\frac{H \vdash A: \diamond \quad \Gamma; A \vdash H: \diamond}{\Gamma \vdash H; A: \diamond} \quad \text{(WF-STATE)}
\end{array}$$

Figure 9: Well-formed states

$$\Gamma, l: \tau; \mathcal{R}; \mathcal{Q} \vdash \mathbf{join} \ l: (\tau, \mathcal{R}, \mathcal{Q}) \quad \text{(T-JOIN)}$$

Figure 10: Typing rules for run-time expressions

phase without becoming quiescent on c . State

$$\{c: \langle _, \{l, l', \dots\}, _ \rangle\}; \{l: (\emptyset, _, _), l': (\{c: _ \}, _, _), \dots\}$$

is also ill formed, since activity l is mentioned as registered with clock c and is not part of l 's local view (which is \emptyset). Any other activity registered with c (l' in the example) is bound to deadlock because l will never quiesce on c .

In order to state our results we must be able to type check run-time expressions as well as machine states. The typing rules for run-time expression **join** l and for machine states and activities are depicted in Figures 10 and 11. The type of a **join** l expression (rule T-JOIN) is that of activity l . Notice that **join** l is the result of evaluating a **finish** e expression (rule R-FINISH) that is, in fact, the type of e (rule T-FINISH). It is worth noticing that a heap H is well typed if each clock is assigned to a different singleton type and if the clocks allocated in the heap are exactly those of typing Γ (rule T-HEAP where \mathbf{C} is the set of all clocks). Moreover, activities may only resume on registered clocks. An activity (V, e, A) has the type of its expression e (rule T-ACT), which must unregister from all its clocks before terminating, since after evaluating e it is expected that the set of registered clocks should be empty. Rule T-STATE incorporates the definition of well-formed states into the type system. The remaining typing rules should be easy to follow.

Lemma 1 (Weakening). *Let a be a variable or a clock name.*

1. If $H \vdash A: \diamond$ then $H, c: h \vdash A: \diamond$.
2. If $\Gamma; \mathcal{R} \vdash v: \tau$ then $\Gamma, a: \tau' \vdash v: \tau$.
3. If $\Gamma \vdash V: \mathcal{R}$ then $\Gamma, a: \tau \vdash V: \mathcal{R}$.
4. If $\Gamma; \mathcal{R}; \mathcal{Q} \vdash e: T$ then $\Gamma, a: \tau; \mathcal{R}; \mathcal{Q} \vdash e: T$.

$$\begin{array}{c}
\frac{\Gamma; \mathcal{R} \vdash c_1 \dots c_n: \mathcal{R}}{\Gamma \vdash \{c_1: _, \dots, c_n: _ \}: \mathcal{R}} \quad \text{(T-VIEW)} \\
\frac{\Gamma \vdash V: \mathcal{R} \quad \mathcal{Q} \subseteq \mathcal{R} \quad \Gamma; \mathcal{R}; \mathcal{Q} \vdash e: (\tau, \emptyset, \emptyset) \quad \Gamma \vdash A}{\Gamma \vdash (V, e, A): \tau} \quad \text{(T-ACT)} \\
\frac{\Gamma \vdash a_1: \tau_1 \quad \dots \quad \Gamma \vdash a_n: \tau_n}{\Gamma, l_1: \tau_1, \dots, l_n: \tau_n \vdash \{l_1: a_1, \dots, l_n: a_n\}} \quad \text{(T-ACT-SET)} \\
\frac{\Gamma; \mathcal{R} \vdash c_1 \dots c_n: \mathcal{R} \quad \{c_1, \dots, c_n\} = \text{dom} \Gamma|_{\mathbf{C}}}{\Gamma \vdash \{c_1: h_1, \dots, c_n: h_n\}} \quad \text{(T-HEAP)} \\
\frac{\Gamma \vdash H; A: \diamond \quad \Gamma \vdash H \quad \Gamma \vdash A}{\Gamma \vdash H; A} \quad \text{(T-STATE)}
\end{array}$$

Figure 11: Typing rules for machine states

- Proof outline.*
1. By induction on the derivation of the typing rules. Case WF-ACT-SET-E is direct. For case WF-ACT-SET we use the induction hypothesis to prove that $H, \{c: h\} \vdash A: \diamond$ and $H, \{c: h\} \vdash A': \diamond$; the remaining conditions are given by the hypotheses.
 2. By inspecting the typing rules.
 3. We apply rule T-VIEW to typify the clocks of the view, then we prove T-CLOCK-SEQ with (2).
 4. By induction on the derivation of the typing relation. Cases T-MAKE, T-NEXT, and T-JOIN are direct. Case T-RESUME and T-DROP are proved similarly, using (2) to typify clock v . The proof for cases T-FINISH, T-ASYNC, and T-LET follow by induction hypothesis. For case T-ASYNC we also use rule T-CLOCK-SEQ and (2) to typify the clocks of the arguments. \square

Notice we do not allow heap weakening for it would introduce in the type environment clocks not present in the state.

Lemma 2 (Substitution). *If $\Gamma; \mathcal{R} \vdash v: \tau$ and $\Gamma, x: \tau; \mathcal{R}; \mathcal{Q} \vdash e: T$ then $\Gamma; \mathcal{R}; \mathcal{Q} \vdash e[v/x]: T$.*

Proof outline. For T-VALUE we analyse two cases: when the value is the variable being substituted, and when it is not replaced. For the former case, we apply Lemma 1 on the first hypothesis. For the latter case we use rule T-VAR. Rules T-MAKE, T-NEXT, and T-JOIN are direct. Cases T-DROP and T-RESUME follow by induction hypothesis. Rule T-ASYNC is the most complex. For the clocks being shared we use Lemma 1 and the second hypothesis. For the expression being spawned we apply the induction hypothesis. Rule T-FINISH is proved similarly to T-ASYNC, but simpler, since T-FINISH has no arguments and its set of clocks is empty. \square

Lemma 3 (Preservation for activities). *If $\Gamma \vdash H$ and $\Gamma \vdash V: \mathcal{Q}$ and $\mathcal{Q} \subseteq \mathcal{R}$ and $\Gamma; \mathcal{R}; \mathcal{Q} \vdash e: T$ and $\Gamma \vdash A$ and $l \in \text{dom} H$ and $H; (V, e, A) \rightarrow_l H'; A''; (V', e', A')$, then $\Gamma' \vdash H'$ and $\Gamma' \vdash V': \mathcal{Q}'$ and $\mathcal{Q}' \subseteq \mathcal{R}'$ and $\Gamma'; \mathcal{R}'; \mathcal{Q}' \vdash e': T$ and $\Gamma' \vdash A', A''$, for some $\Gamma' \supseteq \Gamma$.*

Proof outline. Despite the scary look of the statement, its proof is a routine inspection of the rules in the relation $H; (V, e, A) \rightarrow_l H'; A''; (V', e', A')$. \square

Lemma 4 (Preservation for X10_{|clocks}). *If $\Gamma \vdash S$ and $S \rightarrow S'$ then $\Gamma' \vdash S'$ and $\Gamma \subseteq \Gamma'$.*

Proof outline. By induction on the derivation of the relation $S \rightarrow S'$. In all cases we build the derivation tree for $\Gamma \vdash S$ using rules T-STATE, T-ACT-SET, and T-ACT, collect the hypotheses, use the above Lemmas, and then build a tree for $\Gamma' \vdash S'$ using the same typing rules. The base cases are when the derivation ends with rules R-LET-VAL and R-LET. For R-LET-VAL we take $\Gamma' = \Gamma$ and use the substitution Lemma 2. For R-LET we use the (specially crafted) preservation for activities (Lemma 3), as well as the weakening (Lemma 1) for the extant activity set A . The induction step is when derivation ends with rule R-ACTIVITY; in this case we use the Weakening Lemma. \square

Theorem 5 (Type Safety). *If $\Gamma \vdash S$ and $S \rightarrow^* S'$, then $S' \notin \text{Error}$.*

Proof outline. We first establish that $\Gamma' \vdash S'$ using preservation (Lemma 4). Then we proceed by contradiction. The contradiction is proved by induction on the definition of Error predicate. For the base cases of **resume**, **drop**, and **async** we build the derivation trees for the errors in Figure 5, to conclude that $\Gamma' \vdash V : \mathcal{R}$ and $\Gamma'; \mathcal{R} \vdash c : \mathbf{clock}(\alpha)$. Sequent $\Gamma' \vdash V : \mathcal{R}$ is derived from rule T-VIEW, which effectively establishes a *one-to-one* correspondence between the clock names c in V and the singleton types α in \mathcal{R} . On the other hand, sequent $\Gamma'; \mathcal{R} \vdash c : \mathbf{clock}(\alpha)$ is derived via rule T-WF-C, which says that $\alpha \in \mathcal{R}$. Since $\alpha \in \mathcal{R}$, the correspondence allows us to conclude that $c \in \text{dom} V$. Establishing that $v \in \text{dom} H$ is easier. Given that $\Gamma' \vdash H$, we conclude that $\text{dom} H = \text{dom}(\Gamma' |_{\mathbf{C}})$, and from sequent $\Gamma'; \mathcal{R} \vdash c : \mathbf{clock}(\alpha)$ we know that $c \in \text{dom} \Gamma'$, hence done. \square

In Section 3 we introduced a loading function that builds the initial machine state corresponding to a given expression. Such a state is typable if the expression is. Let $\text{load}(e)$ be defined as the state $\emptyset; \{l : (\emptyset, \mathbf{let } x = e \mathbf{ in unit}, \emptyset)\}$. Then we have:

Lemma 6. *If $\Gamma; \emptyset; \emptyset \vdash e : T$ then $\Gamma \vdash \text{load}(e)$.*

Our final result guarantees that a well-typed expression does not reduce to an error.

Corollary 7. *If $\Gamma; \emptyset; \emptyset \vdash e : T$ then $\text{load}(e)$ does not reduce to an Error.*

Proof. From the lemma above and Theorem 5. \square

We anticipate a *progress property* for typable processes. Typability ensures that processes do not get stuck when dropping a clock that is not in its clock set anymore, or when otherwise trying to access a clock that it not allocated in the heap. The remaining case is **next** where the activity waits for set C_1 (the set of quiescent clocks the activity is registered with) to grow until becoming (together with C_2 —the set of clocks that have already advance their phase) the clock set of the activity. And this is bound to happen for both **next** and **drop**, since in each activity both implicitly resume all clocks. We foresee as well that typability also rules out programs that deadlock, since **finish** expressions can only use clocks created in its body expressions.

6 Discussion and future work

We study two synchronisation constructs of X10: a primitive finish that waits for the termination of activities (lightweight threads), and clocks (a generalisation of barriers). To better understand the language we define an operational semantics and a type system (alternative to the constraint-based system [9]) for a subset of X10 called $\text{X10}|_{\text{clocks}}$. Our main result is type safety for typable programs (Theorem 5).

Our semantics represents clocks in the heap as triples $\langle p, R, Q \rangle$ relying on two sets for recording the registered activities R and the quiesced activities Q on a clock. Implementing operations that work

with sets is costly; for instance rule R-NEXT needs to compute sets C_1 and C_2 , by checking if sets R and Q are equal, and then verify if $C_1 \cup C_2 = \text{dom}V$. Should we make a real life implementation of the proposed semantics, set operations would have a significant impact on performance. We sketch a much faster approach that chooses to represent clocks as triples $\langle p, r, q \rangle$ describing the clock phase, as before, but taking r and q as the cardinal numbers of sets R and Q . With this representation we lose information about the identity of the activities registered with a clock and, in particular, we cannot determine if an activity has already resumed in the current phase (*vide* rules R-ASYNC and R-RESUME). To overcome this problem we need to enrich the clock local view with an indicator of whether an activity has performed a resume in the current phase. Thus, a clock local view becomes a pair $\langle p, b \rangle$ containing the current clock phase p (as before) and the resume boolean indicator b , describing when the activity has resumed. With this information it is straightforward to adapt rules R-ASYNC, R-MAKE, R-RESUME, R-NEXT, and R-DROP. For instance, rule R-RESUME only updates the clock global view ($q \leftarrow q + 1$) whenever its local view indicator is false. Also, rule R-NEXT needs to set r to zero when advancing the clock global phase, and to clear the indicator b upon advancing the clock local phase. Checking that all activities registered with a clock have quiesced amounts to compare two integer values ($r = s$), instead of two sets R and Q as before. The main reasons for not adopting the semantics just sketched are that the chosen semantics needs fewer rules and is easier to read and understand.

We intend to investigate imperative features of the language, specially those related with clocks, and also other language constructs. The finish construct is not only used to wait for the termination of sub-activities, but also, as the language reference reads, “A collection point for uncaught exceptions generated during the execution of S [the body of a finish]” [8, page 196]. Enriching our model with exceptions seems like a natural, promising follow-up of our work. The language report also reads, “X10 does not contain a register statement that would allow an activity to discover a clock in a data structure and register itself on it” (page 208); we would like to study type-safe extensions to the language that might alleviate this restriction in controlled situations. Furthermore, we expect to extend our results to $X10|_{\text{clocks}}$ equipped with recursion or some form of iteration. Futures are a form of a function that evaluates asynchronously, like an activity, but can be *forced* to finish locally to return a value. The semantics of a future, in what regards termination, is like the **finish** construct, but its use cases are different. We would also like to allow futures to register themselves with clocks, a feature missing in X10.

Phasers are a coordination construct that unifies collective and point-to-point synchronisations with performance results competitive to existing barrier implementations [10]. Phasers can be seen as an extension over clocks that allow for more fine-grained control over synchronisation modes. *Phaser accumulators* are reduction constructs for dynamic parallelism that integrate with phasers [11]. Although further investigation is needed, we believe our work can be extended to accommodate phasers and phaser accumulators, specially with regards to the operational similarities between clocks and phasers.

Acknowledgements

The authors would like thank the anonymous referees for constructive criticisms and detailed comments.

References

- [1] Shivali Agarwal, Rajkishore Barik, Vivek Sarkar & Rudrapatna K. Shyamasundar (2007): *May-happen-in-parallel analysis of X10 programs*. In: *Proceedings of PPOPP'07*, ACM, pp. 183–193,

- doi:10.1145/1229428.1229471.
- [2] Philippe Charles, Christian Grothoff, Vijay Saraswat, Christopher Donawa, Allan Kielstra, Kemal Ebcioglu, Christoph von Praun & Vivek Sarkar (2005): *X10: an object-oriented approach to non-uniform cluster computing*. In: *Proceedings of OOPSLA'05*, ACM, pp. 519–538, doi:10.1145/1094811.1094852.
 - [3] Frederica Damera, David A. George, V. Alan Norton & Gregory F. Pfister (1988): *A Single-Program-Multiple-Data computational model for EPEX/FORTRAN*. *Parallel Computing* 7(1), pp. 11–24, doi:10.1016/0167-8191(88)90094-4.
 - [4] Matteo Frigo, Charles E. Leiserson & Keith H. Randall (1998): *The implementation of the Cilk-5 multi-threaded language*. In: *Proceedings of PLDI'98*, ACM, pp. 212–223, doi:10.1145/277650.277725.
 - [5] Rajiv Gupta (1989): *The fuzzy barrier: a mechanism for high speed synchronization of processors*. *SIGARCH Computer Architecture News* 17(2), pp. 54–63, doi:10.1145/68182.68187.
 - [6] Doug Lea (2000): *A Java fork/join framework*. In: *Proceedings of JAVA'00*, ACM, pp. 36–43, doi:10.1145/337449.337465.
 - [7] Jonathan K. Lee & Jens Palsberg (2010): *Featherweight X10: a core calculus for async-finish parallelism*. In: *Proceedings of PPOPP'10*, ACM, pp. 25–36, doi:10.1145/1693453.1693459.
 - [8] Vijay Saraswat (2011): *Report on the Programming Language X10, version 2.12*. Technical Report, IBM Research.
 - [9] Vijay Saraswat & Radha Jagadeesan (2005): *Concurrent clustered programming*. In: *Proceedings of CONCUR'05*, LNCS 3653, Springer, pp. 353–367, doi:10.1007/11539452_28.
 - [10] Jun Shirako, David M. Peixotto, Vivek Sarkar & William N. Scherer (2008): *Phasers: a unified deadlock-free construct for collective and point-to-point synchronization*. In: *Proceedings of ICS'08*, ACM, pp. 277–288, doi:10.1145/1375527.1375568.
 - [11] Jun Shirako, David M. Peixotto, Vivek Sarkar & William N. Scherer (2009): *Phaser accumulators: A new reduction construct for dynamic parallelism*. In: *Proceedings of IPDPS'09*, IEEE Computer Society, pp. 1–12, doi:10.1109/IPDPS.2009.5161071.
 - [12] Chau-Wen Tseng (1995): *Compiler optimizations for eliminating barrier synchronization*. In: *Proceedings of PPOPP'95*, ACM, pp. 144–155, doi:10.1145/209936.209952.
 - [13] Vasco T. Vasconcelos, Francisco Martins & Tiago Cogumbreiro (2010): *Type Inference for Deadlock Detection in a Multithreaded Typed Assembly Language*. In: *Post-proceedings of PLACES'09, EPTCS* 17, pp. 95–109, doi:10.4204/EPTCS.17.8.