

A Featherweight Model for Chorded Languages

Alexis Petrounias Sophia Drossopoulou Susan Eisenbach

January 15, 2009

Abstract

Chords are a concurrency mechanism of object-oriented languages inspired by the *join* of the Join-Calculus. We present SCHOOL, the Small Chorded Object-Oriented Language, a featherweight model which aims to capture the essence of the concurrent behaviours of chords. Our model serves as a generalisation of chorded behaviours found in existing experimental languages such as Polyphonic C^\sharp . Furthermore, we study the interaction of chords with fields by extending SCHOOL to include fields, resulting in f SCHOOL. Fields are orthogonal to chords in terms of concurrent behaviours. We show that adding fields to SCHOOL does not change its expressiveness by means of an encoding between the two languages.

1 Introduction

The *chord* construct is a concurrency mechanism inspired by the *join* from the Join-Calculus [13]. Chords were implemented in Polyphonic C^\sharp [5], an extension of C^\sharp , and are also available in C_ω [3]. Their proponents say that their use will raise the level of abstraction concurrent programs are written in, and make the development of correct programs easier.

Furthermore, the inclusion of chords in C_ω means they will be used in conjunction with imperative concurrency constructs such as monitors and threads. In order to study the interactions between other language constructs and chords, it is necessary that we fully understand the behaviour of chords themselves.

We therefore provide a featherweight model of chorded languages, namely the Small Chorded Object-Oriented Language (SCHOOL). Our formalisation aims at capturing the essence of the concurrent behaviour of chords, and thus lacks many common object-oriented features which are not directly related to chords (such as fields and exceptions).

SCHOOL turned out to be surprisingly small, featuring just four rules and permutation. Yet, the model can describe a rich set of chorded behaviours, including generalisations of those found in current experimental languages such as Polyphonic C^\sharp or Scala with *join* [9].

We also examine the interaction of chords with fields, a standard feature of many object-oriented languages; we find that fields are orthogonal to chords in terms of concurrent program behaviours, and can be encoded using only chords.

We extend SCHOOL by adding fields, resulting in f SCHOOL, and define an encoding between the two languages. We then prove equivalence between encoded programs and hence show that fields do not add expressive power to SCHOOL. Proofs are available from: slurp.doc.ic.ac.uk/chords/school/. Our work has evolved out of a substantially more complicated model [11], called SCHOOL₀ in this paper, which we have worked on previously.

The structure of this paper is as follows: section 2 provides an overview of chords and programming with them; section 3 presents the formalisation of SCHOOL, several of its properties including soundness and progress, and a comparison with the older model for SCHOOL and with Polyphonic C[#]; section 4 presents f SCHOOL and its properties, the encoding of f SCHOOL in SCHOOL, and the proof of equivalence between SCHOOL and encoded f SCHOOL programs; section 6 overviews related work; and section 7 concludes.

2 Overview of Chords

Chorded programs consist of classes which define chords. A chord consists of a header and a body. The header consists of at most one *synchronous* method signature and zero or more *asynchronous* method signatures, while the body consists of the expressions to be executed.

The body of a chord executes when an object has received an invocation for each of the method (signatures) in its header. In general, multiple invocations are required to execute the body. The simultaneous presence of invocations for each of the methods reflects the *join* notion. Hence a chord header can be seen as a guard for the execution of the body.

When a join occurs the participating asynchronous methods' invocations are consumed, and their arguments are passed to the body of the chord. When multiple invocations of the same method are present there is a non-deterministic choice as to which invocation is consumed.

A method can appear at most once in any given chord header, however, methods can participate in multiple chord headers. If multiple chords can join by consuming the same method invocation, then the choice of which chord joins is unspecified.

The invocation of a synchronous method results in the invoking thread *blocking* until a suitable join occurs. Again, there may be a choice of which chord will join and unblock the thread if the method participates in more than one chord. Once the join occurs the invoker thread is unblocked and executes the chord body, potentially resulting in a return value.

Asynchronous methods return immediately, and their return type must be *async*, a subtype of *void*, as there is no body of expressions to result in a return value. A chord which does not contain a synchronous method is called asynchronous; such chords will not have an invoking thread blocking on them. Therefore, when an invocation for each of their participating asynchronous methods is present, the chord can join and its body is executed in a new thread.

The following chord implements an unbounded buffer:

```
Object get( ) & async put( Object o ) { return o; }
```

Invoking `get`, which is synchronous, will result in the invoker blocking until a value (of type *Object* in this case) is returned. The body of the chord will execute only when the chord joins, which requires an invocation of the asynchronous method `put` to be present.

When a join occurs, the invocation of `put` is consumed. Multiple invocations of `put` will result in the invocations being placed into the execution for joining later with invocations of `get` (the invocations are “queued”).

Example: a Countdown Latch

Consider the following problem: an expression executing in a thread T requires the results of several other threads R_1, \dots, R_n , before it can continue execution. These threads complete asynchronously, and hence T must wait for their completion (the order of completion is unimportant). One solution is to use a countdown latch; we compare a chorded implementation of such a latch (listing 1) with a traditional implementation (in Java, listing 2).

A countdown latch can be expressed as a single chord. The thread T invokes the `await` method and passes the value of `permits`, i.e. the number of other threads it wishes to wait for. Since this is a synchronous call, it will block T until the chord joins.

The threads R_1, \dots, R_n , upon completion of their tasks, will signal the latch by invoking `countDown`. This invocation is asynchronous, as these threads should not have to wait for T to perform any operations.

The first time `countDown` is invoked, it will cause the chord to join, and its body will execute. The value of `permits` will be decremented by one, and `await` will be synchronously invoked again with the new value. After n joins, the value of `permits` will be zero, and the recursion will unwind, effectively unblocking T . Since invocations to `countDown` are queued, they can arrive during the execution of the chord body by T without affecting the end result.

In a Java version of the countdown latch, a field, `permits`, holds the number of other tasks the thread T must wait for. The information contained in this field is shared between the two methods (and hence between threads T and R_1, \dots, R_n , and hence must be encapsulated in the latch object instead of remaining local to a method, as in the case of the chorded implementation.

The method `await` sets the initial value of `permits` and then blocks, which is achieved by invoking `wait`, available to all objects in Java. This method will return upon a `notify` invocation on the same object. Because the field `permits` is shared between multiple threads, it needs to be accessed within synchronised blocks to avoid race conditions.

As a thread R_i completes and invokes `countDown`, the value of `permits` is decremented. If the value reaches zero, it means that thread R_i is in fact the last thread to notify of completion, and hence thread T must be unblocked (via invocation of `notify`).

```

1 class CountdownLatch {
2     void await(int permits) & async countDown() {
3         if (--permits > 0) { await(permits); }
4     }
5 }

```

Listing 1: Chorded Countdown Latch

```

1 class CountdownLatch {
2     int permits;
3     void await(int permits) throws InterruptedException {
4         synchronized(this) {
5             this.permits = permits;
6             while (permits > 0) { wait(); }
7         }
8     }
9     void countDown() {
10        synchronized(this) {
11            if (--permits == 0) { notify(); }
12        }
13    }
14 }

```

Listing 2: Java Countdown Latch

In Java it is possible for the `wait` method to *spuriously* wake up and return. Thus, we must place the call to `wait` inside the loop `while (permits > 0)`, which ensures the latch condition is honoured.

Finally, the semantics of `wait` are such that it must be invoked within a *synchronized* block, but upon invocation it releases the monitor acquired by this block. Then, upon waking up, it attempts to re-acquire the monitor (so there can be no race-condition in the window of time between `wait` returning and the evaluation of the loop condition, which accesses `permits`).

Comparing the two implementations, we can identify the following problematic aspects of the Java solution:

- Shared variables: can lead to race-conditions; they must have all access protected by means of mutual exclusion.
- Magic library methods: reliance on methods outside the programming language in order to perform primitive operations such as concurrency.
- Spurious notifications: programs have to include code in order to deal with memory model deficiencies, increasing complexity and reducing comprehensibility.
- Irrelevant delays: threads R_1, \dots, R_n must wait for the internal logic of the countdown latch to execute every time they invoke `countDown`. The logic of the latch, however, is irrelevant to these threads.

- Invisible semantics: the interaction between `wait` and *synchronized* blocks is invisible in terms of language constructs.

The above are not specific to Java; programs in C^\sharp suffer from similar problems.

3 SCHOOL

SCHOOL is a small object-oriented language. The constructs of the language are limited to classes which define chords, and object instantiations of these classes which reside in a heap. Classes exist within a simple, single-inheritance hierarchy, and methods and chords can be overridden.

The generalised chord derivation of the previous section closely resembles the chords of Polyphonic C^\sharp , which features the *async* return type, a special subtype of *void*, to indicate asynchronous methods; this return type is not, however, necessary in order to implement asynchronous methods once we observe that all methods of return type *void* can be executed asynchronously, as the invoker is not expecting a return value.

The chords of SCHOOL are similar to those presented in section 2, however, SCHOOL does not feature the *async* type; instead, a method which has a return type of *void* can be invoked either synchronously or asynchronously, depending on whether this method is found in the synchronous or asynchronous part of a chord header. SCHOOL programs have potentially more valid behaviours than chords with explicit *async* return types; hence, SCHOOL constitutes a more general model of chored behaviours than existing experimental languages.

Furthermore we require exactly one argument for each method, as parameter passing is not particularly interesting in the context of chords. We name the argument of method m as $m.x$. The value of the last expression evaluated in a body becomes the return value of the chord.

3.1 Abstract Syntax and Program Representation

Figure 1 provides an overview of syntax and program representation. The syntax of SCHOOL expressions, *Expr*, consists of method calls $e.m(e)$, variables x , object creation $new\ c$, the special receiver *this*, and the *null* value.

We use Id^m for the set of method names, Id^c for the set of class names, and Id^x for the set of variable names. SCHOOL programs are represented using tuples of mappings. Programs consist of three mappings: inheritance, method signatures, and chords.

The first component maps a class name to another class name: $Id^c \rightarrow Id^c$, expressing the direct superclass relationship between the classes; the superclass of a class c is $P\downarrow_1(c)$.

The second component maps a class name c and a method name m to the method's signature: $Id^c \times Id^m \rightarrow MethSig$; in order to look up the method signature for method m of class c we use $P\downarrow_2(c, m)$. The signature itself consists of a return type, a name, and the class type of the single argument: $t\ m(c)$.

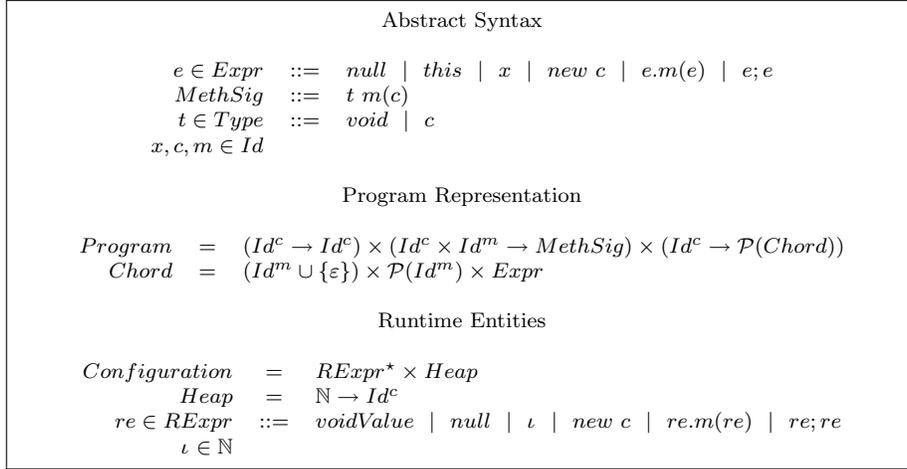


Figure 1: SCHOOL overview.

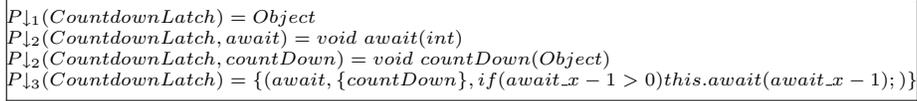


Figure 2: SCHOOL representation of the Countdown Latch program.

The third component maps a class name to the set of chords defined in the class: $Id^c \rightarrow \mathcal{P}(Chord)$; to obtain the chords of class c we use $P\downarrow_3(c)$. We encode chords as triplets $(Id^m \cup \{\varepsilon\}) \times \mathcal{P}(Id^m) \times Expr$. The first element is either the name of a synchronous method, or the symbol ε (indicating that the chord lacks a synchronous method and hence is asynchronous). The second element is a set of asynchronous method names. The third is the body of the chord.

The program of listing 1 is represented in abstract syntax as in figure 2, ignoring the rule about number of parameters, and assuming that arithmetic and integers are defined in the language.

3.2 Execution

Execution of SCHOOL expressions is described by a term rewriting system in which a configuration, consisting of a collection of expressions, e_i , and a heap, h , evaluate into a new configuration:

$$e_1, \dots, e_n, h \longrightarrow e'_1, \dots, e'_m, h'$$

where the heap is a mapping of object addresses to their class names:

$$h : \mathbb{N} \rightarrow Id^c$$

Evaluation Contexts	
$E[\cdot] ::= [\cdot] \mid E[\cdot].m(e) \mid \iota.m(E[\cdot]) \mid E[\cdot]; e$	
Evaluation Rules	
$\frac{h(\iota) \text{ is undefined}}{E[\text{new } c], h \longrightarrow E[\iota], h[\iota \mapsto c]}$	$\frac{}{E[z; e], h \longrightarrow E[e], h}$ SEQ
$\frac{h(\iota)=c \quad m \in \bigcup_{\chi \in P\downarrow_3(c)} \chi \downarrow_2 \quad E[\cdot] \neq [\cdot]}{E[\iota.m(v)], h \longrightarrow E[\text{voidValue}], \iota.m(v), h}$ ASYNC	
$\frac{h(\iota)=c \quad (m, \{m_1, \dots, m_k\}, e) \in P\downarrow_3(c)}{E[\iota.m(v)], E_1[\iota.m_1(v_1)], \dots, E_k[\iota.m_k(v_k)], h \longrightarrow E[e[\iota/\text{this}, v/m_x, v_1/m_{1-x}, \dots, v_k/m_{k-x}], E_1[\text{voidValue}_1], \dots, E_k[\text{voidValue}_k]], h}$ JOIN	
$\frac{h(\iota)=c \quad (\varepsilon, \{m_1, \dots, m_k\}, e) \in P\downarrow_3(c)}{E_1[\iota.m_1(v_1)], \dots, E_k[\iota.m_k(v_k)], h \longrightarrow E_1[\text{voidValue}_1], \dots, E_k[\text{voidValue}_k], E[e[\iota/\text{this}, v_1/m_{1-x}, \dots, v_k/m_{k-x}], h]}$ STRUNG	
$\frac{\bar{e} \cong \bar{e}''\bar{e}'''' \quad \bar{e}'', h \longrightarrow \bar{e}''', h' \quad \bar{e}'''\bar{e}'''' \cong \bar{e}'}{\bar{e}, h \longrightarrow \bar{e}', h'}$ PERM	

Figure 3: SCHOOL operational semantics.

and the number of expressions may change from n to m , as new expressions are *spawned* when asynchronous chords join and their bodies execute. Furthermore, threads never terminate in the sense that ground expressions are not removed from the execution. Hence it is always the case that $m \geq n$.

We also use the shorthand \bar{e} for several, concurrent expressions, and thus we also have:

$$\bar{e}, h \longrightarrow \bar{e}', h'$$

Since \bar{e} denotes concurrency not sequentiality, expressions in \bar{e} are separated by “,” instead of “;”.

We describe SCHOOL using a structural operational semantics found in figure 3; we use the variable v to designate acceptable values for arguments to method calls, which consist of all expressions other than *voidValue*, and the variable z to designate all irreducible values (*null*, ι , *voidValue*). The evaluation rules are:

- NEW: creates a new object of a given class and allocates a previously undefined heap address which now maps to the object; the result is the new address.
- SEQ: discards an irreducible value and enables the evaluation of the next expression in a sequential composition; the final value in a sequential com-

position cannot be discarded, and this allows the final value of a chord’s body to become the return value of the chord’s synchronous method.

- **ASYNC**: places an asynchronous invocation of a method (of *void* type and appearing in at least one asynchronous part of a chord header) into the execution, available for joining later. The invocation immediately returns *voidValue*. The condition $E[\cdot] \neq [\cdot]$, requiring the evaluation context to not be empty, is necessary so that we avoid infinite reductions of the form:

$$\begin{aligned} \iota.m(v), h &\longrightarrow \text{voidValue}, \iota.m(v), h &\longrightarrow \\ \text{voidValue}, \text{voidValue}, \iota.m(v), h &\longrightarrow \dots \end{aligned}$$

- **JOIN**: selects a chord in which a synchronous method participates and *joins* this chord by consuming the corresponding asynchronous invocations (and replacing each by *voidValue*). The actual arguments are mapped to the formal arguments of the chord’s body, which becomes the new evaluating expression.
- **STRUNG**: selects an asynchronous chord which is *strung*, i.e can join. Similar to the JOIN rule, all the asynchronous invocations are consumed, and actual arguments are mapped to the formal arguments of the chord’s body, which will execute concurrently with the rest of the expressions.
- **PERM** enables the non-deterministic selection of expressions to evaluate and the reordering of expressions in the execution. The notation $\bar{e} \cong \bar{e}'$ means that \bar{e}' is a permutation of \bar{e} .

The selection of which *strung* chord to join happens at two levels: multiple receiver objects may have asynchronous invocations enabling the joining of a chord, and an object may feature multiple chords which currently can join.

For example, consider two threads using the countdown latch of figure 2: one thread invokes `await(2)` and another thread invokes `countDown()` twice. We assume that the heap h has an instantiated countdown latch at address ι , and we assume standard operational semantics for sequential execution and if-statements; an execution sequence is in figure 4. Because this example requires a conditional rule for the if-statement, which is not part of SCHOOL as we have focused on the concurrent aspect of chords, we show this rule in figure 5 for the sake of the example; another option would be to encode the if-statement using methods. We will not deal with the commonly employed *else* part, and we will assume booleans have been defined in the language.

On line 1 we start with the two threads. On line 2 the first invocation of `countDown` from the second thread is evaluated through the ASYNC rule and results in the invocation being replaced with *voidValue* and the invocation itself being placed into the execution as a new thread. On line 3 the first thread’s invocation of `await` joins with the newly available invocation: the body of the chord is the currently evaluating expression for the first thread and *voidValue* is the value of the third thread. On line 4 the *voidValue* of the second thread

1		$\iota.await(2), \iota.countDown(null); \iota.countDown(null), h$
2	$\xrightarrow{\text{ASYNC}}$	$\iota.await(2), voidValue; \iota.countDown(null), \iota.countDown(null), h$
3	$\xrightarrow{\text{JOIN}}$	$(if\ 1 > 0)\ \iota.await(1), voidValue; \iota.countDown(null), voidValue, h$
4	$\xrightarrow{\text{SEQ}}$	$(if\ 1 > 0)\ \iota.await(1), \iota.countDown(null), voidValue, h$
5	$\xrightarrow{\text{IF}}$	$\iota.await(1), \iota.countDown(null), voidValue, h$
6	$\xrightarrow{\text{ASYNC}}$	$\iota.await(1), voidValue, voidValue, \iota.countDown(null), h$
7	$\xrightarrow{\text{JOIN}}$	$(if\ 0 > 0)\ \iota.await(0), voidValue, voidValue, voidValue, h$
8	$\xrightarrow{\text{IF}}$	$voidValue, voidValue, voidValue, voidValue, h$

Figure 4: Example of the Countdown Latch execution in SCHOOL.

Abstract Syntax	Evaluation Contexts
$e \in Expr ::= \dots \mid if\ (e)\ e$	$E[\cdot] ::= \dots \mid if\ (E[\cdot])\ e$
Evaluation Rule	
$\frac{e' = \begin{cases} e & \text{if } v = true \\ voidValue & \text{otherwise} \end{cases}}{E[if\ (v)\ e], h \longrightarrow E[e'], h} \text{ IF}$	

Figure 5: SCHOOL if-statements.

is removed since it is sequential composition, and on line 5 the condition in the if-statement is evaluated. In the next few lines similar executions take place, but on line 8 the condition is false and the chord executing in the first thread eventually returns.

The heaps of SCHOOL are particularly simple: they only record the existence of objects and their associated type, and serve only as a means of storing objects and allowing method invocations to determine their receivers. This is possible because SCHOOL lacks many of the object oriented features that require a more structured heap. In particular, queues of messages to asynchronous methods are implicit and are determined by examination of the configuration. If this were not possible, then explicit queues of method invocations would have to be encoded, per object, in the heap, and the language semantics would be significantly more complex.

3.3 Types, Well-Formedness and Subject Reduction

The type system of SCHOOL (presented in figure 6) is nominal, with class names constituting types. There is an implicit empty top superclass called *Object*, and all other class types are defined through class declarations (we write $P \vdash_{cl} c$ to state that class c is declared in program P). In addition to classes, there is a primitive type, *void*. We write $P \vdash_{tp} t$ to indicate that t is a type in program P . Terms in SCHOOL are typed using a context, Γ , and a heap, h . We use the same type system for source level and runtime expressions.

Type Definitions	
$\frac{P \downarrow_1(c) \text{ is defined}}{P \vdash_{\delta_{cl}} c} \text{ DEF-CLASS}$	$\frac{}{P \vdash_{\delta_{cl}} \textit{Object}} \text{ DEF-CLASS-OBJECT}$
$\frac{P \vdash_{\delta_{cl}} c}{P \vdash_{\delta_{tp}} c} \text{ DEF-TYPE}$	$\frac{}{P \vdash_{\delta_{tp}} \textit{void}} \text{ DEF-TYPE}$
Typing Judgements	
$\frac{}{P, \Gamma, h \vdash \textit{voidValue} : \textit{void}} \text{ T-VOID}$	$\frac{x \in \{\textit{this}\} \cup Id^x}{P, \Gamma, h \vdash x : \Gamma(x)} \text{ T-THISX}$
$\frac{h(\iota) = c}{P, \Gamma, h \vdash \iota : c} \text{ T-ADDRESS}$	$\frac{P \vdash_{\delta_{cl}} c}{P, \Gamma, h \vdash \textit{new } c : c} \text{ T-NEW}$
$\frac{}{P, \Gamma, h \vdash \textit{null} : c} \text{ T-NULL}$	$\frac{P, \Gamma, h \vdash e_1 : c \quad P, \Gamma, h \vdash e_2 : t \quad P \downarrow_2(c, m) = t_r \quad m(t)}{P, \Gamma, h \vdash e_1.m(e_2) : t_r} \text{ T-INV}$
$\frac{P, \Gamma, h \vdash e_0 : _ \quad P, \Gamma, h \vdash e : t}{P, \Gamma, h \vdash e_0; e : t} \text{ T-SEQ}$	

Figure 6: SCHOOL type system.

For a class c of a program P to be well-formed, denoted $P \vdash c$ (see definition 1), it is necessary to fulfil four criteria:

1. It must have a superclass in P .
2. All method signatures appearing in its superclass must also appear in the class.
3. Each method signature must participate in at least one chord; if the method's return type is *not void* then the method signature must participate as the synchronous part of a chord.
4. For each chord appearing in the class, all of its participating methods must be present; furthermore, the asynchronous methods (if any) must have a return type of *void*; finally, the type of a chord body must agree with the return type of the synchronous method, or be *void* if the chord is asynchronous.

Definition 1 (SCHOOL Well-Formed Classes)

$P \vdash c$ iff :

1. $P \vdash_{\circ_{ct}} P \downarrow_1(c)$
2. $P \downarrow_2(P \downarrow_1(c), m)$ defined $\implies P \downarrow_2(c, m) = P \downarrow_2(P \downarrow_1(c), m)$
3. $P \downarrow_2(c, m) = t \ m(_)$ $\implies \begin{cases} m \in \bigcup_{\chi \in P \downarrow_3(c)} \{\chi \downarrow_1\} \cup \chi \downarrow_2 & \text{if } t = \text{void} \\ m \in \bigcup_{\chi \in P \downarrow_3(c)} \{\chi \downarrow_1\} & \text{otherwise} \end{cases}$
4. $(\alpha, \{m_1, \dots, m_n\}, e) \in P \downarrow_3(c) \implies$
 $\forall i \in 1..n : \exists t_i : P \downarrow_2(c, m_i) = \text{void} \ m_i(t_i)$
 $P, m_1 _x \mapsto t_1, \dots, m_n _x \mapsto t_n, \text{this} \mapsto c, \gamma, _ \vdash e : t$
 where, either $\alpha = \varepsilon, \gamma = \varepsilon, t = \text{void}$
 or $\alpha = m, P \downarrow_2(c, m) = t \ m(t'), \gamma = m _x \mapsto t'$

A program P is well-formed, if all declared classes are well-formed.

Definition 2 (SCHOOL Well-Formed Programs)

$\vdash P$ iff $P \vdash_{\circ_{ct}} c \implies P \vdash c$

Subject reduction (theorem 1) guarantees for a well-formed program P , a heap h , and well-typed expressions \bar{e} , execution results in a heap h' and well-typed expressions \bar{e}' ; furthermore, the type of each expression remains unchanged, and if the evaluation spawned a new thread, then this thread is well-typed too (with type void). The theorem relies on appropriate substitution of variables during method invocation (definition 3) and three auxiliary lemmas: substitution preserves types (lemma 1), evaluation preserves the types of objects in the heap (lemma 2), and the evaluation of a single expression preserves its type (lemma 3).

Definition 3 (Substitution of Variables)

For a substitution $\sigma = \text{Id} \cup \{\text{this}\} \rightarrow \text{Addr}$, a heap h , and an environment Γ , we have: $P, \Gamma, h \vdash \sigma$ iff $\forall id \in \text{dom}(\Gamma) : P, \Gamma, h \vdash \sigma(id) : \Gamma(id)$

Lemma 1 (Substitution Preserves Types)

$$\left. \begin{array}{l} P, \Gamma, h \vdash e : t \\ P, \Gamma, h \vdash \sigma \end{array} \right\} \implies P, \Gamma, h \vdash [e]_{\sigma} : t$$

Proof 1 By structural induction on the derivation of $P, \Gamma, h \vdash e : t$; case analysis on the typing judgement used, and use of definition 3. The cases for T-VOID, T-NULL, T-NEW, and T-ADDRESS are immediate as they do not involve variables. The cases for T-SEQ, T-SUBCLASS, and T-INV are satisfied inductively. Finally, in the case for T-THISX we observe from the first premise of the lemma, $P, \Gamma, h \vdash e : t$, and the conclusion of the typing judgement that e is of the form x and hence $P, \Gamma, h \vdash x : \Gamma(x)$, so $P, \Gamma, h \vdash e : \Gamma(x)$ where $t = \Gamma(x)$; we also observe from the second premise of the lemma, $P, \Gamma, h \vdash \sigma$, and the definition of σ , that $P, \Gamma, h \vdash \sigma(id) : \Gamma(id)$, and therefore $P, \Gamma, h \vdash \sigma(x) : \Gamma(x)$; we conclude that $P, \Gamma, h \vdash [e]_{\sigma} : \Gamma(x)$, which satisfies the lemma.

Lemma 2 (Evaluation Preserves Types of Objects in the Heap)

$$e, h \longrightarrow e', h' \implies \forall \iota \in \text{dom}(h) : h(\iota) = h'(\iota)$$

Proof 2 By case analysis on the evaluation rule used in the derivation of $e, h \longrightarrow e', h'$. The only applicable rules are NEW and SEQ. For the case of NEW we obtain from its premise that $h(\iota)$ is undefined, while from its conclusions we obtain $h[\iota \mapsto c]$; we therefore conclude that $\forall \iota \in \text{dom}(h) : h(\iota) = h'(\iota)$, which satisfies the lemma. For the case of SEQ we obtain that $h = h'$, which immediately satisfies the lemma.

Lemma 3 (Evaluation of Single Expression Preserves its Type)

$$\left. \begin{array}{l} \vdash P \\ P, \Gamma, h \vdash e : t \\ e, h \longrightarrow e', h' \end{array} \right\} \implies P, \Gamma, h' \vdash e' : t$$

Proof 3 By structural induction on the derivation of $e, h \longrightarrow e', h'$; case analysis on the last evaluation rule applied. The only applicable rules are NEW and SEQ.

For the case of NEW we obtain from its premise that $h(\iota)$ is undefined, while from its conclusions we obtain that the form of e is new c , and that of e' is ι , as well as that the resulting heap is $h' = h[\iota \mapsto c]$; from the second premise of the lemma and the form of e we can apply the typing judgement T-NEW (the first premise of the lemma satisfies the premise of this typing judgement), and obtain $P, \Gamma, h \vdash \text{new } c : c$, and hence $t = c$; from the form of e' and the resulting heap h' we can apply the typing judgement T-ADDRESS and conclude that $P, \Gamma, h' \vdash \iota : c$, which satisfies the lemma.

For the case of SEQ we obtain that the form of e is $z; e_0$ for some expression e_0 , that the form of e' is e_0 , and that $h = h'$; from the second premise of this lemma and the form of e we can apply the typing judgement T-SEQ and obtain $P, \Gamma, h \vdash e : t'$ so $t = t'$; from the form of e' and the fact that the heap remains unchanged we conclude that $P, \Gamma, h' \vdash e' : t'$, which satisfies the lemma.

Theorem 1 (SCHOOL Subject Reduction)

$$\left. \begin{array}{l} \vdash P \\ e_1, \dots, e_n, h \longrightarrow e'_1, \dots, e'_n, e, h' \\ \forall i \in 1..n : P, \Gamma, h \vdash e_i : t_i \end{array} \right\} \implies \begin{array}{l} \forall i \in 1..n : P, \Gamma, h' \vdash e'_i : t_i \\ \wedge P, \Gamma, h' \vdash e : \text{void if } e \neq \varepsilon \end{array}$$

Proof 4 By structural induction on the derivation of $e_1, \dots, e_n, h \longrightarrow e'_1, \dots, e'_n, e, h'$; case analysis on the last evaluation rule applied, and use of definition 3 and lemmas 1, 2, and 3. In the following cases we will use the shorthand notion $\bar{e} \equiv e_1, \dots, e_n$ and $\bar{e}' \equiv e'_1, \dots, e'_n, e$.

—

The case for PERM is satisfied inductively, where we appeal to lemma 3 in the case of a single expression evaluating (through NEW or SEQ), or the inductive hypothesis of this theorem in all other cases (ASYNC, JOIN, and STRUNG).

Furthermore, in each of the three cases the heap remains unchanged, and using the above result that evaluation does not change the types of objects in the heap (lemma 2), we immediately obtain that the type of an expression in the initial configuration remains unchanged in the final configuration when the expression does not evaluate.

—

In the case for ASYNC we observe that the form of \bar{e} is $E[\iota.m(v)]$ and that of \bar{e}' is $E[\text{voidValue}], \iota.m(v)$, while the heap remains unchanged, hence $h' = h$; from the premises we know that the receiver object is of type c , since $h(\iota)=c$, and we know that the method m participates in the asynchronous part of at least one chord, since $m \in \bigcup_{\chi \in P \downarrow_3(c)} \chi \downarrow_2$.

Since \bar{e} is well-typed (by the second premise of the theorem), from the form of \bar{e} we can apply the typing judgement T-INV and obtain (a) $P, \Gamma, h \vdash \iota.m(v) : t_r$ for some type t_r , (b) $P, \Gamma, h \vdash \iota : c'$ for some class c' , (c) $P, \Gamma, h \vdash v : t_v$ for some type t_v , and (d) $P \downarrow_2(c', m) = t_r m(t_v)$.

Since $h(\iota)=c$, from (b) we obtain that $c' = c$, and subsequently from (d) we obtain $P \downarrow_2(c, m) = t_r m(t_v)$.

Since the method m participates asynchronously in at least one chord, and since the program is well-formed (from the first premise of the theorem), we refer to well-formedness and obtain $P \downarrow_2(c, m) = \text{void } m(_)$, which immediately gives us that $t_r = \text{void}$ (t_v from (c) is consistent with this). Using our knowledge of t_r and referring back to (a), we obtain $P, \Gamma, h \vdash \iota.m(v) : \text{void}$; furthermore, using the fact that the heap remains unchanged, we obtain $P, \Gamma, h' \vdash \iota.m(v) : \text{void}$.

From the form of $E[\text{voidValue}]$ and from the fact that the heap remains unchanged, we can apply the typing judgement T-VOID and obtain $P, \Gamma, h' \vdash E[\text{voidValue}] : \text{void}$.

Therefore, the single expression from \bar{e} did not change type, and the new expression created in \bar{e}' is of type void , which satisfies the theorem.

—

In the case for JOIN we observe that the form of \bar{e} is $E[\iota.m(v)], E_1[\iota.m_1(v_1)], \dots, E_k[\iota.m_k(v_k)]$ and that of \bar{e}' is $E[[e]_\sigma], E_1[\text{voidValue}_1], \dots, E_k[\text{voidValue}_k]$, where we have a substitution $\sigma = [\iota/\text{this}, v/m_x, v_1/m_{1-x}, \dots, v_k/m_{k-x}]$, while the heap remains unchanged, hence $h' = h$; from the premises we know that the receiver object is of type c , since $h(\iota)=c$, and we know that the method m participates in the synchronous part of at least one chord, as well as that the asynchronous methods m_1, \dots, m_k form the chord's asynchronous part, since $(m, \{m_1, \dots, m_k\}, e) \in P \downarrow_3(c)$.

From the chord $(m, \{m_1, \dots, m_k\}, e)$ and well-formedness we know that $\exists t, t' : P \downarrow_2(c, m) = t m(t')$ and $\forall i \in 1..k : \exists t_i : P \downarrow_2(c, m_i) = \text{void } m_i(t_i)$, and $P, m_x \mapsto t', m_{1-x} \mapsto t_1, \dots, m_{n-x} \mapsto t_k, \text{this} \mapsto c, _ \vdash e : t$. From the typing judgement T-ADDRESS and knowing $h(\iota)=c$, we obtain $P, \Gamma, h \vdash \iota : c$

(and consequently $P, \Gamma, h \vdash \text{this} : c$ through T-THISX); and then from the typing judgement T-INV we obtain $P, \Gamma, h \vdash \iota.m(v) : t \wedge P, \Gamma, h \vdash v : t'$ and $\forall i \in 1..k : P, \Gamma, h \vdash \iota.m_i(v_i) : \text{void} \wedge P, \Gamma, h \vdash v_i : t_i$.

From the above we notice that for all identifiers in the substitution σ the type of the original identifier matches the type of the substituted identifier, and hence by definition 3 we conclude that $P, \Gamma, h \vdash \sigma$, and therefore by lemma 1 and the fact that the heap remains unchanged we obtain $P, \Gamma, h' \vdash [e]_\sigma : t$.

From the typing judgement T-VOID and the form of e'_1, \dots, e'_k we obtain $\forall i \in 1..k : P, \Gamma, h' \vdash \text{voidValue}_i : \text{void}$.

Therefore, all expressions have retained their type, which satisfies the theorem.

—

In the case for STRUNG we observe that the form of \bar{e} is $E_1[\iota.m_1(v_1)], \dots, E_k[\iota.m_k(v_k)]$ and that of \bar{e}' is $E_1[\text{voidValue}_1], \dots, E_k[\text{voidValue}_k], E[[e]_\sigma]$, where we have a substitution $\sigma = [\iota/\text{this}, v_1/m_1.x, \dots, v_k/m_k.x]$, while the heap remains unchanged, hence $h' = h$; from the premises we know that the receiver object is of type c , since $h(\iota) = c$, and we know that the asynchronous methods m_1, \dots, m_k form a chord's asynchronous part, since $(\varepsilon, \{m_1, \dots, m_k\}, e) \in P \downarrow_3(c)$.

From the chord $(\varepsilon, \{m_1, \dots, m_k\}, e)$ and well-formedness we know that $\forall i \in 1..k : \exists t_i : P \downarrow_2(c, m_i) = \text{void } m_i(t_i)$, and $P, m_1.x \mapsto t_1, \dots, m_n.x \mapsto t_k, \text{this} \mapsto c, _ \vdash e : \text{void}$. From the typing judgement T-ADDRESS and knowing $h(\iota) = c$, we obtain $P, \Gamma, h \vdash \iota : c$ (and consequently $P, \Gamma, h \vdash \text{this} : c$ through T-THISX); and then from the typing judgement T-INV we obtain and $\forall i \in 1..k : P, \Gamma, h \vdash \iota.m_i(v_i) : \text{void} \wedge P, \Gamma, h \vdash v_i : t_i$.

From the above we notice that for all identifiers in the substitution σ the type of the original identifier matches the type of the substituted identifier, and hence by definition 3 we conclude that $P, \Gamma, h \vdash \sigma$, and therefore by lemma 1 and the fact that the heap remains unchanged we obtain $P, \Gamma, h' \vdash [e]_\sigma : \text{void}$.

From the typing judgement T-VOID and the form of e'_1, \dots, e'_k we obtain $\forall i \in 1..k : P, \Gamma, h' \vdash \text{voidValue}_i : \text{void}$.

Therefore, all expressions have retained their type, and the new expression created in \bar{e}' is of type void , which satisfies the theorem.

Because we have focused on the concurrent aspect of chorded languages, we have not dealt with inheritance in SCHOOL beyond a rudimentary subtype relation for classes. Well-formedness of classes requires that all method signatures appearing in a super class also appear in a sub class, which means that a given class will contain definitions for all the methods appearing in its entire class hierarchy; furthermore, since all defined methods must appear in at least one chord, all sub classes must either “copy” the chords of their super class, or “override” those chords (as long as all methods are used). The consequence for subject reduction is that all method and chord look-up operations always succeed for well-formed programs, which greatly simplifies the type system.

The heaps of SCHOOL have a minimal impact on properties of the language such as subject reduction; specifically, the only evaluation rule which modifies the heap is NEW, which results in a monotonically increasing heap where no previous object is affected. The typing judgement T-ADDRESS is the only judgement which involves the heap and determines the type of an object in order to look-up methods and chords of its class. The rest of the evaluation rules leave the heap unchanged, which greatly simplifies the above proof as all expressions which do not change in the current evaluation step can immediately be shown to retain their type.

Some care must be taken when performing substitution during chord joining in SCHOOL; in contrast to traditional languages, where a method signature contains all the information necessary for substitution upon invocation, in SCHOOL a synchronous method can participate in potentially multiple chords which consist of different patterns of asynchronous methods. Hence, the same synchronous method may receive different substitutions depending on which chord it joins. The fourth part of well-formedness for classes ensures that the appropriate substitution takes place and the proof of subject reduction appeals to this construction in order to show that all valid chord look-ups satisfy the constraints of substitution.

3.4 Progress

Progress for SCHOOL (see theorem 2 below) is stated as follows: either a configuration can evaluate through application of an evaluation rule, or no evaluation rule is applicable and the configuration is then either *terminated* or *blocked*. Termination occurs when all expressions are *ground values* (or *null-pointers*, see below), while a blocked configuration is one in which at least one expression is blocked, and no evaluation rule is applicable (the rest of the expressions are either blocked or grounded (or null-pointers)).

Blocked Expressions

An expression is *blocked* (see definition 4 below) if it is not a ground value or null-pointer and it is not possible to apply an evaluation rule to further evaluate it. In SCHOOL, the only such form of expression is an invocation on a non-null receiver. If the method invocation is asynchronous and in a non-empty evaluation context, it can always be further evaluated through the ASYNC rule and hence is never blocked (part 1 of the definition). If the method participates synchronously in a chord, then at least one asynchronous method invocation necessary for the joining of that chord must be missing from the configuration (part 2 of the definition). If the method participates asynchronously in a purely asynchronous chord, then at least one *other* asynchronous method invocation necessary for that chord must be missing from the configuration (part 3 of the definition). If the method participates asynchronously in a synchronous chord it is never considered blocked and is not captured by this definition (we will

consider the invocation a ground value instead if the chord cannot join, see below).

Definition 4 (SCHOOL Blocked Expressions)

For a program P , a configuration \bar{e}, h , an expression $e \in \bar{e}$ is blocked iff it is of the form $\iota.m(_)$ where $h(\iota)=c$ and additionally and the following three hold:

1. $m \in \bigcup_{\chi \in P \downarrow_3(c)} \chi \downarrow_2 \implies E[\cdot] = [\cdot]$
2. $\forall (m, \{\bar{m}\}, _) \in P \downarrow_3(c), \bar{m} \neq \emptyset :$
 $\exists m' \in \bar{m} : \iota.m'(_) \notin \bar{e} \setminus e$
3. $\forall (\varepsilon, \{\bar{m}\}, _) \in P \downarrow_3(c), m \in \bar{m}, \bar{m} \setminus m \neq \emptyset :$
 $\exists m' \in \bar{m} : \iota.m'(_) \notin \bar{e}$

From the above definition we notice that if a method participates asynchronously in any chord and its context is not empty, then the evaluation rule ASYNC is immediately applicable and thus the configuration is not blocked (hence part 1 requires $E[\cdot] = [\cdot]$).

Furthermore, if a chord consists of a single method, either synchronous (a degenerate chord) or asynchronous (a trivial chord), then it does not cause the configuration to be blocked, as we can immediately apply either the JOIN rule or the STRUNG rule, respectively (hence parts 2 and 3 require $\bar{m} \neq \emptyset$ and $\bar{m} \setminus m \neq \emptyset$ respectively).

Finally, if a method participates both as the synchronous and as an (or the) asynchronous method of a chord, at least *two* invocations of the method must be present (required in order to apply the JOIN rule), otherwise the configuration is blocked (hence part 2 requires the missing invocation to be from the set $\bar{e} \setminus e$).

Because parts 2 and 3 quantify over *all* chords, an invocation will indeed only be considered blocked when no evaluation rule can be applied to further evaluate it.

Ground Values

An expression is a *ground value* (see definition 5 below) when its form is either *voidValue*, or ι , or *null*, as there is no rule through which it can further evaluate. However, it is not obvious whether an asynchronous method invocation, $\iota.m(v)$, placed in the configuration through the ASYNC rule, and thus having an empty context, $E[\cdot] = [\cdot]$, constitutes a ground value or is considered blocked when it cannot participate in any evaluation.

This problem can be stated as the question: is a configuration, where at least one chord's join is partially satisfied (by the presence of a subset of method invocations), blocked? It turns out that the answer is a matter of preference, as both options lead to the same theorem of progress (the proof of the theorem, however, *is* affected by the choice of answer).

We have decided to consider irreducible asynchronous method invocations as ground values when they can only partially satisfy the join of a synchronous

chord, and blocked when they can only partially satisfy the join of an asynchronous chord. The justification for this choice is that in the case of synchronous chords, emphasis is placed on the synchronous method, and therefore it is the absence of a synchronous invocation which results in a blocked execution; on the other hand, for asynchronous chords, all participating methods are equally important, and hence the absence of any invocation renders the configuration blocked.

Therefore, an asynchronous invocation is considered a *ground value* when it is an asynchronous invocation with an empty context, so that it cannot immediately evaluate through ASYNC (part 1 of the definition), it does not participate as a synchronous method of any chord, as then it would be considered blocked not ground (part 2 of the definition), it does not participate in any purely asynchronous chords, as again it would be considered blocked not ground (part 3 of the definition), and finally whenever it participates in the asynchronous part of a synchronous chord, it is considered ground only when the chord is partially satisfied: either an invocation of the synchronous method of the chord must be missing, and/or an invocation of some other asynchronous method of the chord must be missing (part 4 of the definition).

Definition 5 (SCHOOL Ground Values)

For a program P and a configuration \bar{e}, h , an expression $e \in \bar{e}$ is a *ground value* iff it is of the form *voidValue*, or ι , or *null*; or of the form $\iota.m(_)$ where $h(\iota)=c$ and additionally the following four hold:

1. $m \in \bigcup_{\chi \in P \downarrow_3(c)} \chi \downarrow_2 \implies E[\cdot] = [\cdot]$
2. $\bar{A}(m, _, _) \in P \downarrow_3(c)$
3. $\bar{A}(\varepsilon, \bar{m}, _) \in P \downarrow_3(c) : m \in \bar{m}$
4. $\forall (m', \bar{m}, _) \in P \downarrow_3(c), m' \neq \varepsilon, m \in \bar{m} :$
 $\iota.m'(_) \notin \bar{e} \quad \wedge \quad \exists m'' \in \bar{m} : \iota.m''(_) \notin \bar{e}$

Because part 4 quantifies over *all* synchronous chords in which the method participates asynchronously, the invocation will be considered ground only when no synchronous chord in which it participates is fully satisfied. Therefore, if the asynchronous invocation fully satisfies one chord and partially satisfies another, it is not considered ground.

Null-Pointers

Considering the case of a method invocation on a *null* address, of the form $null.m(_)$; we notice that this expression can type check, however, it cannot reduce as there is no evaluation rule which accepts *null* in place of an ι . Therefore, we consider such invocations as *null-pointer expressions*, and in terms of progress they are treated as ground values. Whether a configuration with a null-pointer expression can be considered terminated (as would a configuration consisting of only ground values) is subject to interpretation of the meaning of termination one wishes to give.

Progress

Using the above definitions can show progress for all well-typed SCHOOL configurations: a well-typed configuration can either be further evaluated, or it consists of ground values (or null-pointers) only (has terminated), or it is blocked (at least one blocked expression when no evaluation rule is applicable and the rest of the expressions are either blocked or ground (or null-pointers)).

Theorem 2 (SCHOOL Progress)

$$\left. \begin{array}{l} \vdash P \\ P, _, h \vdash \bar{e} : \bar{t} \end{array} \right\} \implies \left\{ \begin{array}{l} \exists \bar{e}', h' : \bar{e}, h \longrightarrow \bar{e}', h' \\ \vee \bar{e} \text{ consists of ground values or} \\ \quad \text{null-pointer expressions only} \\ \vee \bar{e}, h \text{ is blocked} \end{array} \right.$$

Proof 5 *By case analysis on the applicability of evaluation rules; when no evaluation rule can be applied we perform a case analysis on the participation of asynchronous methods in chords and use definitions 4 and 5 to determine whether they are ground values or whether they are blocked.*

3.5 Comparison Between SCHOOL and Polyphonic C[#]

There are several structural differences between Polyphonic C[#] chords and SCHOOL chords:

- While Polyphonic C[#] has methods and chords, we unify the two and treat Polyphonic C[#] methods as SCHOOL chords with a single synchronous method in their header.
- We have no *async* type, and allow a method of return type *void* to appear in synchronous and asynchronous parts of chord headers.
- In Polyphonic C[#] asynchronous methods can only be overridden by asynchronous methods, and synchronous methods can only be overridden by synchronous methods; SCHOOL does not have this restriction.
- Similarly, in Polyphonic C[#] methods are either synchronous or asynchronous; SCHOOL does not have this restriction either.
- Polyphonic C[#] class hierarchies are acyclic, while this is not a constraint imposed by SCHOOL.

The last three Polyphonic C[#] constraints are not necessary for type soundness, and so have not been included in SCHOOL.

3.6 Simplifications from Previous Model

Our original formulation of SCHOOL is presented in [11]; we will refer to it as SCHOOL₀. SCHOOL is substantially simpler than SCHOOL₀; we give a brief overview of the simplifications here.

In SCHOOL₀ method invocation is modelled as message passing between objects: the argument of the method being invoked is sent to the receiving object and subsequently placed within a queue corresponding to that method. Hence the structure of objects consists of their class name and a set of queues, one for each asynchronous method.

In the case of a SCHOOL₀ synchronous method call, waiting for a join to occur on a chord of a particular receiver object, it suffices to inspect only that object's queues. However, in the case of asynchronous chords there is no method call waiting for a join. Hence, it is necessary to inspect all objects whose class contains at least one asynchronous chord. Thus, detection of deadlock is a property of the expressions and the heap.

In SCHOOL deadlock can now be detected by inspecting only the set of executing expressions. This enables us to treat synchronous and asynchronous chords in a uniform way with respect to deadlock detection.

We have removed the special type *async*, using *void* as the return type of asynchronous methods. This makes the type system and proof of subject reduction simpler, as it dispenses with converting *async* to *void*.

Additionally, the use of the *void* return type for asynchronous methods enables us non-deterministically to chose whether such methods are invoked synchronously or asynchronously: all methods of return type *void* can participate in both the synchronous and asynchronous parts of chords.

We have also removed those constructs which can be programmatically encoded: null pointers, exceptions, and sequencing.

The semantics of SCHOOL is much simpler than that of SCHOOL₀: there are fewer rules (from nine versus five), and the rules are shorter.

3.7 Comparison Between SCHOOL and Polyphonic C[#]

There are several structural differences between Polyphonic C[#] chords and SCHOOL chords:

- While Polyphonic C[#] has methods and chords, we unify the two and treat Polyphonic C[#] methods as SCHOOL chords with a single synchronous method in their header.
- We have no *async* type, and allow a method of return type *void* to appear in synchronous and asynchronous parts of chord headers.
- In Polyphonic C[#] asynchronous methods can only be overridden by asynchronous methods, and synchronous methods can only be overridden by synchronous methods; SCHOOL does not have this restriction.
- Similarly, in Polyphonic C[#] methods are either synchronous or asynchronous; SCHOOL does not have this restriction either.
- Polyphonic C[#] class hierarchies are acyclic, while this is not a constraint imposed by SCHOOL.

```

1 class CountdownLatch {
2   int permits;
3   void await( ) & async countdown( ) {
4     permits = permits - 1;
5     if (permits > 0) { await( ); }
6   }
7 }

```

Listing 3: Chorded Countdown Latch with fields.

The last three Polyphonic C^\sharp constraints are not necessary for type soundness, and so have not been included in SCHOOL.

3.8 Implementation

We have implemented SCHOOL in the form of a virtual machine named Harp, similar to a Java virtual machine (programs are compiled into classes with interpretable bytecodes). The source code for Harp and its compiler, along with example input programs which include the unbounded buffer and the countdown latch can be obtained from the web at the following site:

slurp.doc.ic.ac.uk/chords/harp/

The original design of the compiler and virtual machine, along with a prototype implementation, was by the author; subsequent refinement and new implementations were by Volanakis [19] and Nicolaou [16] as part of final year undergraduate projects supervised by the author in the Department of Computing, Imperial College London, United Kingdom.

4 SCHOOL With Fields: f SCHOOL

In order to focus on chords, in SCHOOL we have omitted fields. However, the functionality of fields can be obtained using only chords, and hence adding fields to SCHOOL would not increase its expressivity. To show this, we extend SCHOOL with field declarations, obtaining f SCHOOL, and then define an encoding that maps f SCHOOL programs into SCHOOL programs, and show that this encoding preserves all observable behaviours. For the remainder of the paper we use the annotation $f \dots$ to denote f SCHOOL entities.

Listing 3 shows the Chorded Countdown Latch example using fields. In the original example (listing 1), the value of `permits` was passed in (recursive) calls to `await()`. Now, we just store its value directly in a field, which a client must initialise before calling `await()`.

Hence, a class in f SCHOOL contains chord definitions as in SCHOOL, plus zero or more field declarations. Expressions in chord bodies can read from and write to an object's fields using the dot operator.

4.1 Abstract Syntax and Program Representation

In order to obtain ${}^f\text{SCHOOL}$ we extend SCHOOL in the following ways:

1. Programs are extended with a component of signature $Id^c \times Id^f \rightarrow c$, which maps the name of a class and a field to the field's type; to look up the type of a field we use ${}^fP\downarrow_4(c, f)$.
2. Objects are extended to contain a mapping of field names to values. The definition of the heap is accordingly extended.
3. To cater for fields, two rules are included which read from and write to a field. The rule for object creation is also extended to initialise the field-to-value mapping of a new object with *null* values.
4. The type system is extended with two judgements for field reads and writes and a judgement for *null* values.
5. The definition of a well-formed class is extended to require that any field declared in a class is inherited and not overridden in a subclass (definition 6).
6. A new definition for well-formed heaps is introduced to ensure that the contents of fields are as according to their static types (definition 8).

The program of listing 3 is represented as in figure 8; it is similar to the original representation from figure 2 of section 3.1, but now includes a fourth element for the field `permits`.

4.2 Types, Well-Formedness and Subject Reduction

We extend the notions of well-formedness to ${}^f\text{SCHOOL}$. Since class and program well-formedness in ${}^f\text{SCHOOL}$ *replaces* the requirements from SCHOOL , we use the symbol \vdash_f to disambiguate. We do not need to do the same for execution (figure 9) and types (figure 10), because the ${}^f\text{SCHOOL}$ evaluation rules and typing judgments *extend* those of SCHOOL .

An ${}^f\text{SCHOOL}$ class c is well-formed, denoted ${}^fP \vdash_f c$, if it well-formed in SCHOOL , and another two additional requirements are met: first, if a field has as a type a class c' , then c' must be declared; second, the type of a field in a class must match the type of the field in the superclass:

Definition 6 (${}^f\text{SCHOOL}$ Additional Well-Formed Classes)

- ${}^fP \vdash_f c$ iff:
1. $P \vdash c$
 2. ${}^fP\downarrow_4(c, f) = c' \implies {}^fP \vdash_{\sigma_{c'}} c'$
 3. ${}^fP\downarrow_4({}^fP\downarrow_1(c), f) = t \implies {}^fP\downarrow_4(c, f) = t$

As in SCHOOL , an ${}^f\text{SCHOOL}$ program is well-formed, denoted $\vdash_f {}^fP$, if all declared classes are well-formed:

Abstract Syntax	
$f_e \in Expr$	$::= null \mid this \mid x \mid new\ c \mid f_e.m(f_e) \mid f_e;f_e \mid f_e.f \mid f_e.f = f_e$
$MethSig$	$::= t\ m(c)$
$t \in Type$	$::= void \mid c$
$x, c, m, f \in Id$	
Program Representation	
$Program$	$= (Id^c \rightarrow Id^c) \times (Id^c \times Id^m \rightarrow MethSig) \times (Id^c \rightarrow \mathcal{P}(Chord))$ $\times (Id^c \times Id^f \rightarrow c)$
$Chord$	$= (Id^m \cup \{\varepsilon\}) \times \mathcal{P}(Id^m) \times Expr$
Runtime Entities	
$Configuration$	$= REExpr^* \times Heap$
$Heap$	$= \mathbb{N} \rightarrow Id^c \times Fields$
$Fields$	$= Id^f \rightarrow \mathbb{N}$
$f_{re} \in REExpr$	$::= voidValue \mid null \mid \iota \mid new\ c \mid f_{re}.m(f_{re}) \mid f_{re};f_{re}$ $\mid f_{re}.f \mid f_{re}.f = f_{re}$
$\iota \in \mathbb{N}$	

Figure 7: f SCHOOL overview.

$fP_{\downarrow 1}(CountdownLatch) = Object$
$fP_{\downarrow 2}(CountdownLatch, await) = void\ await(int)$
$fP_{\downarrow 2}(CountdownLatch, countDown) = void\ countDown(Object)$
$fP_{\downarrow 3}(CountdownLatch) = \{(await, \{countDown\}, this.permits = this.permits - 1;$ $if(this.permits > 0)\ this.await(null);)\}$
$fP_{\downarrow 4}(CountdownLatch, permits) = int$

Figure 8: f SCHOOL representation of the Countdown Latch program.

Definition 7 (f SCHOOL Well-Formed Programs)

$$\vdash_f fP \text{ iff } fP \vdash_{\diamond_{cl}} c \implies fP \vdash_f c$$

An f SCHOOL heap f_h is well-formed, denoted $fP \vdash_f f_h$, if each field of each object can be typed:

Definition 8 (f SCHOOL Well-Formed Heaps)

$$fP \vdash_f f_h \text{ iff } \forall \iota \in f_h : \left. \begin{array}{l} f_h(\iota) = (c, fs) \\ P \vdash_{\diamond_{cl}} c \\ fP_{\downarrow 4}(c, f) = c' \end{array} \right\} \implies fP, _, f_h \vdash_f fs(f) : c'$$

We can show that f SCHOOL is sound by proving subject reduction; the theorem and proof are similar to those for SCHOOL. Because f SCHOOL expressions may contain field accesses, and a field's value may be an address, an expression may evaluate to an address not initially present. Therefore, we need to also take care of the objects' field values, and hence we also require $fP \vdash_f f_h$.

Evaluation Contexts	
$E[\cdot]$	$::= [\cdot] \mid E[\cdot].m(f_e) \mid \iota.m(E[\cdot]) \mid E[\cdot]; f_e \mid E[\cdot].f \mid E[\cdot].f = f_e \mid \iota.f = E[\cdot]$
Evaluation Rules	
$E[\text{new } c], f_h$	$\xrightarrow{f_h(\iota) \text{ is undefined} \quad fP_{\downarrow 4}(c) = \{f_1 _, \dots, f_k _ \}} E[\iota], f_h[\iota \mapsto (c, \{f_1 \mapsto \text{null}, \dots, f_k \mapsto \text{null}\})]$ NEW
$E[z; f_e], f_h$	$\xrightarrow{\text{SEQ}} E[f_e], f_h$
$E[\iota.m(v)], f_h$	$\xrightarrow{f_h(\iota) = (c, _) \quad m \in \bigcup_{\chi \in fP_{\downarrow 3}(c)} \chi \downarrow 2 \quad E[\cdot] \neq [\cdot]} E[\text{voidValue}], \iota.m(v), f_h$ ASYNC
$E[\iota.m_1(v_1), \dots, E_k[\iota.m_k(v_k)], f_h]$	$\xrightarrow{f_h(\iota) = (c, _) \quad (m, \{m_1, \dots, m_k\}, f_e) \in fP_{\downarrow 3}(c)} E[f_e[\iota/\text{this}, v/m_x, v_1/m_{1-x}, \dots, v_k/m_{k-x}], E_1[\text{voidValue}_1], \dots, E_k[\text{voidValue}_k], f_h]$ JOIN
$E_1[\iota.m_1(v_1), \dots, E_k[\iota.m_k(v_k)], f_h]$	$\xrightarrow{f_h(\iota) = (c, _) \quad (\varepsilon, \{m_1, \dots, m_k\}, f_e) \in fP_{\downarrow 3}(c)} E_1[\text{voidValue}_1], \dots, E_k[\text{voidValue}_k], E[f_e[\iota/\text{this}, v_1/m_{1-x}, \dots, v_k/m_{k-x}], f_h]$ STRUNG
$E[\iota.f], f_h$	$\xrightarrow{f_h(\iota) = (c, \{ \dots, f \mapsto v, \dots \})} E[v], f_h$ READ
$E[\iota.f = v], f_h$	$\xrightarrow{f_h(\iota) = (c, fs)} E[v], f_h[\iota \mapsto (c, fs[f \mapsto v])]$ WRITE
$\overline{f_e} \cong \overline{f_e'' f_e''''} \quad \overline{f_e''}, f_h$	$\xrightarrow{\overline{f_e''}, f_h' \quad \overline{f_e'' f_e''''} \cong \overline{f_e'}} \overline{f_e}, f_h$ PERM

Figure 9: f SCHOOL operational semantics.

Theorem 3 (f SCHOOL Subject Reduction)

$$\left. \begin{array}{l} \vdash^{fP} \\ \overline{fP} \vdash^{f_h} \\ fP, \Gamma, f_h \vdash \overline{f_e} : \bar{t} \\ \overline{f_e}, f_h \longrightarrow \overline{f_e'}, f_h' \end{array} \right\} \implies \left\{ \begin{array}{l} fP \vdash^{f_h'} \wedge \\ \exists \bar{t}' \cong \bar{t}, t'' \text{ where } t'' = \text{void or } t'' = \varepsilon : \\ fP, \Gamma, f_h' \vdash \overline{f_e'} : \bar{t}' \end{array} \right.$$

Proof 6 We first prove a theorem for a single expression, similar to theorem 1 for SCHOOL, and then use that in combination with structural induction over the derivation of $\overline{f_e}, f_h \longrightarrow \overline{f_e'}, f_h'$; we perform a case analysis of the last evaluation rule applied.

Type Definitions	
$\frac{fP \downarrow_1(c) \text{ is defined}}{fP \vDash_{\bar{\sigma}_{cl}} c} \text{ DEF-CLASS}$	$\frac{}{fP \vDash_{\bar{\sigma}_{cl}} \text{Object}} \text{ DEF-CLASS-OBJECT}$
$\frac{fP \vDash_{\bar{\sigma}_{cl}} c}{fP \vDash_{\bar{\sigma}_{tp}} c} \text{ DEF-TYPE}$	$\frac{}{fP \vDash_{\bar{\sigma}_{tp}} \text{void}} \text{ DEF-TYPE}$
Typing Judgements	
$\frac{}{fP, \Gamma, fh \vDash \text{voidValue} : \text{void}} \text{ T-VOID}$	$\frac{x \in \{\text{this}\} \cup Id^x}{fP, \Gamma, fh \vDash x : \Gamma(x)} \text{ T-THISX}$
$\frac{fh(\iota) = (c, _)}{fP, \Gamma, fh \vDash \iota : c} \text{ T-ADDRESS}$	$\frac{fP \vDash_{\bar{\sigma}_{cl}} c}{fP, \Gamma, fh \vDash \text{new } c : c} \text{ T-NEW}$
$\frac{}{fP, \Gamma, fh \vDash \text{null} : c} \text{ T-NULL}$	$\frac{fP, \Gamma, fh \vDash f_{e_1} : c \quad fP, \Gamma, fh \vDash f_{e_2} : t \quad fP \downarrow_2(c, m) = t_r \ m(t)}{fP, \Gamma, fh \vDash f_{e_1}.m(f_{e_2}) : t_r} \text{ T-INV}$
$\frac{fP, \Gamma, fh \vDash f_e : c \quad fP \downarrow_4(c, f) = t}{fP, \Gamma, fh \vDash f_e.f : t} \text{ T-READ}$	$\frac{fP, \Gamma, fh \vDash f_{e_1} : c \quad fP, \Gamma, fh \vDash f_{e_2} : t \quad fP \downarrow_4(c, f) = t}{fP, \Gamma, fh \vDash f_{e_1}.f = f_{e_2} : t} \text{ T-WRITE}$
$\frac{fP, \Gamma, fh \vDash f_{e_0} : _ \quad fP, \Gamma, fh \vDash f_e : t}{fP, \Gamma, fh \vDash f_{e_0}; f_e : t} \text{ T-SEQ}$	

Figure 10: f SCHOOL type system.

4.3 Progress

Progress for f SCHOOL can be demonstrated in the same manner as for SCHOOL. The two new evaluation rules, READ and WRITE, ensure that all field reading and writing expressions can evaluate, and hence never contribute towards blocking the configuration; furthermore, the resulting values from reading from and writing to a field will be discarded through the SEQ rule in the event of sequential composition, which ensures progress in such a case.

Theorem 4 (f SCHOOL Progress)

$$\left. \begin{array}{l} \vDash fP \\ fP \vDash fh \\ fP, _, fh \vDash \bar{f}_e : \bar{t} \end{array} \right\} \implies \left\{ \begin{array}{l} \exists \bar{f}_{e'}, \bar{f}_{h'} : \bar{f}_e, fh \longrightarrow \bar{f}_{e'}, \bar{f}_{h'} \\ \vee \bar{f}_e \text{ consists of ground values or} \\ \quad \text{null-pointer expressions only} \\ \vee \bar{f}_e, fh \text{ is blocked} \end{array} \right.$$

```

1 class CountdownLatch {
2   CountdownLatch init_CountdownLatch() {
3     permits(0); return this;
4   }
5   int get_permits() & async permits(int r) {
6     permits(r); return r;
7   }
8   int set_permits(int s) & async permits(int r) {
9     permits(s); return s;
10  }
11  void await() & async countdown() {
12    set_permits(get_permits() - 1);
13    if (get_permits() > 0) { await(); }
14  }
15 }

```

Listing 4: Encoded chorded Countdown Latch with fields.

Proof 7 *Similar to that of theorem 2 with two new straightforward cases for the evaluation rules READ and WRITE. An extra requirement in the case of f SCHOOL is that fields defined in a superclass must also be defined in a subclass, but this is guaranteed by well-formed programs.*

5 Encoding f SCHOOL into SCHOOL

We obtain an encoding of f SCHOOL programs into SCHOOL by employing chords in order to record the current value of each field, read and write values of fields, and to initialise each field with a default value upon object creation.

Listing 4 demonstrates the translation of the chorded countdown latch with fields. There are four chords in the translated class: the first chord, `init_CountdownLatch`, describes object initialization: we create a new countdown latch through `new CountdownLatch().init_CountdownLatch()`. The second chord, with synchronous part `get_permits`, describes reading of field `permits`; through `o.get_permits()` we read the field `permits` from the object `o`. The third chord, with synchronous part `set_permits(int r)`, describes writing field `permits`; through `o.set_permits(v)` we set the field `permits` of the object `o` to the value `v`. The fourth chord, with synchronous part `await`, corresponds to that from listing 3, where field reads/writes have been replaced by calls to `get_permits` and `set_permits`. Furthermore, in listing 4 there is the asynchronous method `permits` whose parameter represents the value of the field `permits`.

Using `new CountdownLatch.init_CountdownLatch()` we create a new countdown latch, and before the new object is returned the invocation of `permits(0)` will be available in the execution.

Figure 11 presents the definition of encoding \mathbb{C} , from f SCHOOL to SCHOOL, for expressions and for programs, where chords encode the behaviours of fields (mapping to values, reading, and writing).

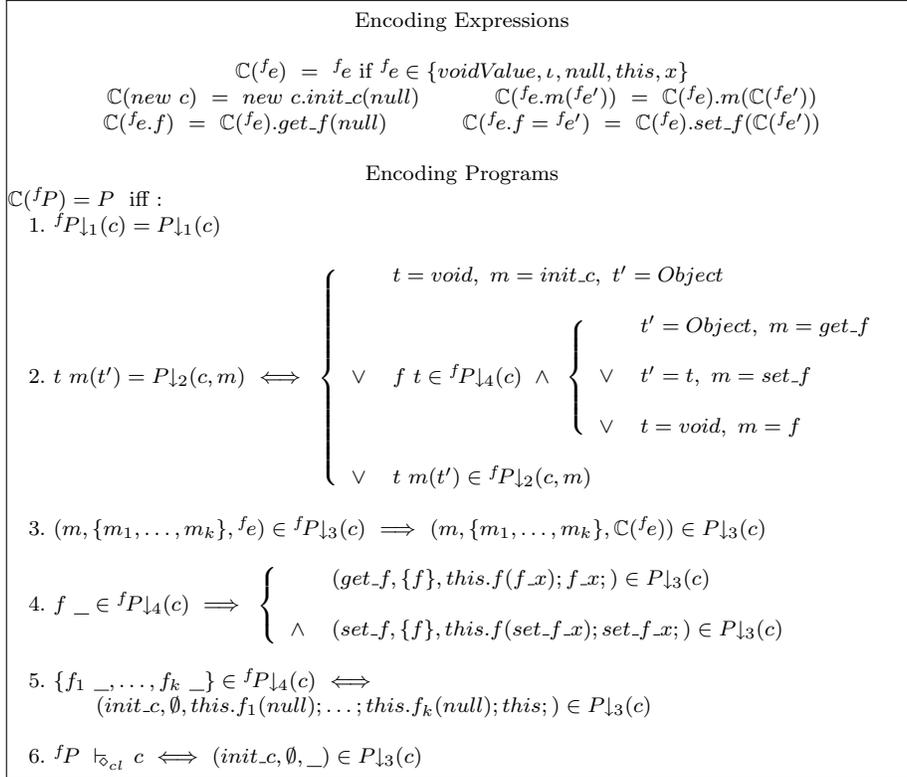


Figure 11: Encoding f SCHOOL into SCHOOL.

For a class c , with a field f of type t , the encoding creates the following method signatures:

$$\begin{aligned} P \downarrow_2(c, f) &= \text{void } f(t) \\ P \downarrow_2(c, \text{get_f}) &= t \ \text{get_f}(\text{Object}) \\ P \downarrow_2(c, \text{set_f}) &= t \ \text{set_f}(t) \\ P \downarrow_2(c, \text{init_c}) &= c \ \text{init_c}(\text{Object}) \end{aligned}$$

The argument type of f is t , to represent the type of the field, and the result type is void , since f is used in the asynchronous parts of the chords. The argument of get_f is discarded, and its type is Object only in order to satisfy our restriction to single-parameter methods. The result type is t . The argument and the result type of set_f is t , to represent the fact that the field has type t , and that the new field value is being returned. The argument type of init_c is Object again for the same reason as it was for get_f , and the result type is t , because the newly initialized object is returned.

The field-reading chord then has the following form:

$$(\text{get_f}, \{f\}, \text{this}.f(f_x); f_x) \in P \downarrow_3(c)$$

The joining of this chord results in the f invocation being consumed, its argument, f_x , eventually being returned as the value of the field. However, as fields should not be consumed by reading, the invocation of f must first be placed back into the execution, via $this.f(f_x)$, so that further reads and writes can take place.

The field writing chord has the following form:

$$(set_f, \{f\}, this.f(set_f_x); set_f_x) \in P\downarrow_3(c)$$

Upon joining, the argument to the set_f method is placed back into the execution (as well as returned). The old value of the field (the argument to the f method) is lost.

From the above two chords we see that the invocation to f acts as a “mutex” to the encoded field; at any given point of execution only one chord instance can have access to the invocation. Only after the chord places the invocation back into the execution can another chord join and consume the invocation. Hence, field reading and writing are atomic, which corresponds with the the atomic evaluation of the f SCHOOL rules READ and WRITE.

Finally, assuming that the class has fields f_1, \dots, f_n , the chord to initialise the fields is the following:

$$(init_c, \emptyset, this.f_1(null); \dots; this.f_n(null); this) \in P\downarrow_3(c)$$

This chord will join when $init_c$ is invoked, and place an $\iota.f_i(null)$ invocation into the execution for each field f_i of the object being initialised, with the default value as argument. Finally, it returns the current object, ι , to its invoker.

We assume that none of the methods in the f SCHOOL program has identifier $get_...$ or $set_...$ or $init_...$, nor the name of a field f from that program.

The return type of the $init_...$ method is dependent on the class name c , as the return type will be different for a subclass of c ; thus, in order for the encoded program to type check, we require a unique name for the initialisation chord’s method (for instance, class D which extends class C will result in two return types for $init_...$, D and C in each class respectively). Were the type system to support covariant overriding of methods, we could use a single name $init$ for all initialisation chords.

Encoding of fields using other language constructs has been studied by amongst others Abadi and Cardelli [1] which uses method invocations and updates in order to obtain the behaviour of fields; this encoding is particularly suitable for functional calculi, as invoking methods and updating their bodies for the remainder of the evaluation does not require assignment to previously defined variables. f SCHOOL, however, focuses on the mutable field construct of many chorded languages and thus does not benefit from a purely functional basis; the encoding of fields for SCHOOL presented above is therefore fundamentally different. Furthermore, the use of asynchronous invocations to represent the current values of fields gives us the framework below for showing correspondence which is not available in the Abadi and Cardelli encoding, as their calculus has no provision for concurrency.

$$\begin{array}{c}
\forall \iota \in \text{dom}(^f h) : ^f h(\iota) = (c, _) \implies h(\iota) = c \\
\bar{e} = \{\iota.f(v) : ^f h(\iota) \downarrow_2(f) = v\} \\
\hline
^f e_1, \dots, ^f e_n, ^f h \simeq \mathbb{C}(^f e_1), \dots, \mathbb{C}(^f e_n), \bar{e}, h \quad \text{C-STRONG} \\
\\
\bar{f}_e \cong \bar{f}_{e'} \quad \bar{f}_{e'}, ^f h \simeq \bar{e}', h \quad \overline{\text{voidValue}} \bar{e}' \cong \bar{e} \\
\hline
\bar{f}_e, ^f h \simeq \bar{e}, h \quad \text{C-STRONG-PERM}
\end{array}$$

Figure 12: Strong Correspondence

5.1 Properties of the Encoding

The goal of the encoding is to show that the behaviour of fields can be obtained using only chords. This essentially means that the encoding satisfied a notion of *soundness* and a notion of *completeness*, where soundness ensures that the behaviours of the original programs (which feature chords and fields) are preserved in the encoded programs (which feature only chords), and completeness ensures that encoded programs do not introduce new behaviours which do not correspond to those of the original programs.

Hence, configurations of original and translated programs must correspond, and execution of both original and encoded programs must lead to respectively corresponding configurations. In the next section we describe *strong correspondence* between an f SCHOOL configuration and its encoding in SCHOOL when expressions, object classes, and field current values are preserved, denoted $\bar{f}_e, ^f h \simeq \bar{e}, h$, with which we show soundness of the encoding.

However, showing completeness is not as straightforward because certain steps of evaluation in f SCHOOL programs require several *intermediate* steps of evaluation, as well as additional concurrent expressions in the encoded SCHOOL program. Field-access chords must join with, and restore, the asynchronous method invocation representing the field’s current value, and object-initialisation chords must populate the configuration with asynchronous invocations consisting of initial field values. Furthermore, due to concurrency, these multiple steps and additional expressions can also be arbitrarily interleaved. Thus, new behaviours can be generated which are not the direct result of the encoding, rather, they depend on non-deterministic interleaving.

Consequently, strong correspondence does not suffice to show completeness, since we need to ensure that although intermediate steps of evaluation “temporarily” break correspondence (such as the current values of fields which are missing immediately after the joining of a field-access chord), eventually, original and encoded programs always further execute and reach strongly corresponding configurations again. Hence, we introduce *weak correspondence*, denoted $\bar{f}_e, ^f h \approx \bar{e}, h$, which allows for these temporary discrepancies during evaluation of intermediate steps.

5.2 Strong Correspondence

In figure 12 we define strong correspondence between f SCHOOL and SCHOOL configurations, in the form of the judgement $\overline{f}e, f_h \simeq \bar{e}, h$.

The first rule, C-STRONG, ensures that an f SCHOOL configuration corresponds with its encoded SCHOOL configuration since it reproduces the behaviours of f SCHOOL expressions and preserves the classes of objects and the current values of fields, by requiring the following:

1. For each expression f_{e_i} in the f SCHOOL configuration, its encoding $\mathbb{C}(f_{e_i})$ appears as an expression in the SCHOOL configuration.
2. All objects in the f SCHOOL heap f_h are also present in the SCHOOL heap h and have the same class.
3. The SCHOOL configuration contains expressions representing the contents of the fields in the f SCHOOL heap f_h .

Note that in the third requirement we use the sequence \bar{e} as a set, when we require that $\bar{e} = \{\dots\}$, and we express that the sequence should contain exactly the elements in this set.

The second rule, C-STRONG-PERM, enables the application of the first rule onto any permutation of the expressions, and allows for any number of additional *voidValue* ground expressions in the SCHOOL configuration (left over from the joining of field-access chords).

Consider as an example the encoding of a program which employes a countdown latch (at address ι_0) where the current value of the `permits` field is 3; the f SCHOOL heap is the following:

$$f_h = \{\iota_0 \mapsto (\text{CountdownLatch}, \{\text{permits} \mapsto 3\})\}$$

and the current configuration is:

$$\iota_0.\text{permits} = 2, f_h$$

The encoded heap in SCHOOL is:

$$h = \{\iota_0 \mapsto \text{CountdownLatch}\}$$

and the encoded configuration in SCHOOL is:

$$\iota_0.\text{set_permits}(2), \iota_0.\text{permits}(3), h$$

By C-STRONG the original program and its encoding indeed correspond:

$$\iota_0.\text{permits} = 2, f_h \simeq \iota_0.\text{set_permits}(2), \iota_0.\text{permits}(3), h$$

If, however, there were two invocations of the `permits` method present in the SCHOOL configuration, the correspondence would not hold:

$$\iota_0.\text{permits} = 2, f_h \not\simeq \iota_0.\text{set_permits}(2), \iota_0.\text{permits}(3), \iota_0.\text{permits}(4), h$$

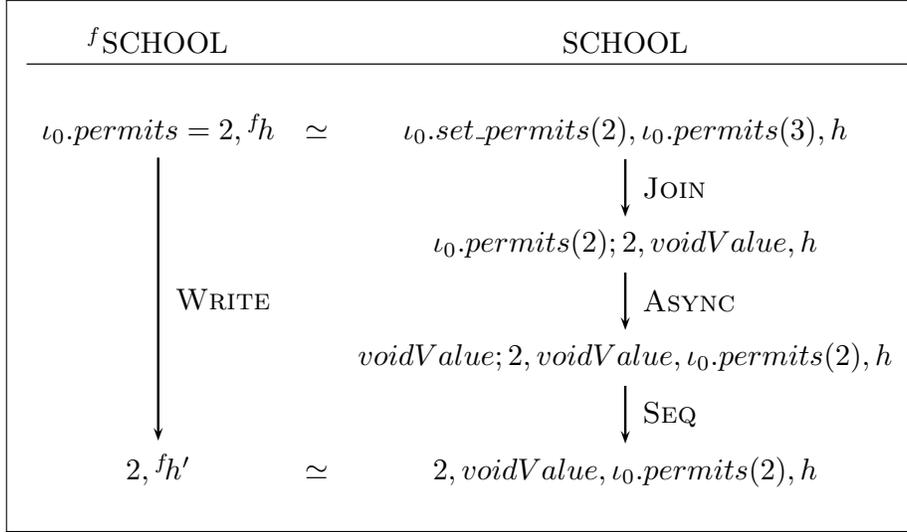


Figure 13: Example of strong correspondence.

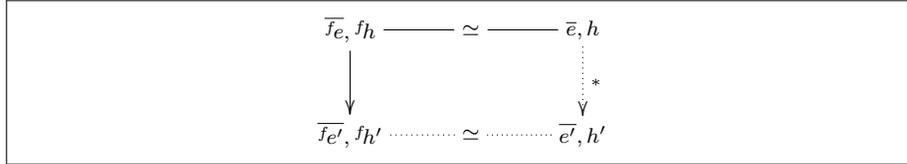


Figure 14: Soundness of the encoding.

Execution of the f SCHOOL program results in a configuration where 2 becomes the current expression, as well as the current value of the `permits` field (through the application of the WRITE rule):

$$\iota_0.\text{permits} = 2, f_h \xrightarrow{\text{WRITE}} 2, f_{h'}$$

where $f_{h'} = f_h[\iota_0 \mapsto (\text{CountdownLatch}, \{\text{permits} \mapsto 2\})]$

The corresponding initial SCHOOL configuration reaches the corresponding final configuration as shown in figure 13.

Using strong correspondence we can show soundness (theorem 5 below); i.e. if the f SCHOOL configuration makes a step, then it is *always possible* to evaluate the strongly corresponding SCHOOL configuration, possibly in multiple steps, so that the final configurations also correspond strongly.

Note that in the theorem below we assume the program fP to be global, in the sense that it is the same program which executes in f SCHOOL on both sides of the implication, and additionally the SCHOOL execution is that of the program $\mathbb{C}(fP)$.

Theorem 5 (Soundness of the Encoding)

$$\left. \begin{array}{l} \overline{f_e, f_h} \simeq \overline{e, h} \\ \overline{f_e, f_h} \longrightarrow \overline{f_{e'}, f_{h'}} \end{array} \right\} \implies \exists \overline{e'}, h' : \overline{e, h} \longrightarrow^* \overline{e'}, h' \quad \wedge \quad \overline{f_{e'}, f_{h'}} \simeq \overline{e'}, h'$$

Proof 8 We first prove a lemma for a single thread by case induction on the execution $f_e, f_h \longrightarrow \overline{f_{e'}, f_{h'}}$; the theorem for many threads then follows.

The diagram from figure 14 illustrates the above theorem; solid lines represent the premises and dotted lines the conclusions.

5.3 Weak Correspondence

In the example from figure 13 we notice that between strongly corresponding configurations, intermediate configurations occur. The aim of weak correspondence is to ensure that although strong correspondence (e.g. field current values) is not preserved during intermediate steps, eventually, we reach configurations which correspond strongly again. Hence, using strong and weak correspondence together, original and encoded program executions correspond throughout.

Under weak correspondence, two field-access chords will never *simultaneously* join for the *same* field of the *same* object. Furthermore, no field-access chord will join before an initial value for the field has been provided by the initialisation chord's body.

Weak correspondence keeps track of which objects and fields are currently being accessed by the SCHOOL expressions of the encoding, and does not allow any other expressions to access them simultaneously. Essentially, the “mutex” behaviour of asynchronous methods representing field current values is enforced.

In figures 15 and 17 we define two forms of the weak correspondence judgement. The first has the form $\varphi \vdash f_e, f_h \approx \overline{e, h}$ and is between an individual f SCHOOL expression and a heap and potentially multiple SCHOOL expressions and a heap. The second has the form $\overline{f_e, f_h} \approx \overline{e, h}$ and is between complete f SCHOOL and SCHOOL configurations, which allows us to apply the individual expression form for each expression in the f SCHOOL configuration. Additionally, we define heap correspondence, which has the form $f_h \approx h$, and which ensures that all objects in an f SCHOOL heap also exist in the SCHOOL heap and have the same class in both heaps.

5.3.1 Individual Expressions

The form of weak correspondence for individual expressions (figure 15) is:

$$\varphi \vdash f_e, f_h \approx \overline{e, h}$$

where the f SCHOOL expression f_e corresponds with potentially multiple SCHOOL expressions \overline{e} , and where an object or field indicated by φ is currently being accessed by the SCHOOL expressions.

When a field f of an object at address ι is being accessed (either by a field-read or field-write chord) the form of φ is $\iota.f$; when an entire object at address ι

$\frac{\varphi \vdash f_e, fh \approx e, \bar{e}, h}{\varphi \vdash f_e, fh \approx z; e, \bar{e}, h} \text{C-WK-VOID}$	$\frac{}{\varepsilon \vdash v, fh \approx v, h} \text{C-WK-VAL}$
$\frac{\varphi \vdash f_e, fh \approx e, \bar{e}, h}{\varphi \vdash f_e.f, fh \approx e.get_f(null), \bar{e}, h} \text{C-WK-READ}$	
$\frac{fh(\iota) = (_, \{ \dots, f \mapsto v, \dots \})}{\iota.f \vdash \iota.f, fh \approx \iota.f(v); v, h} \text{C-WK-READ-BODY}$	
$\frac{\varphi \vdash f_e, fh \approx e, \bar{e}, h \quad \varphi' \vdash f_{e'}, fh \approx e', \bar{e}', h \quad f_e = v \text{ or } \varphi' = \varepsilon}{\varphi \oplus \varphi' \vdash f_e.f = f_{e'}, fh \approx e.set_f(e'), \bar{e}, \bar{e}', h} \text{C-WK-WRITE}$	
$\frac{fh(\iota) \text{ is defined}}{\iota.f \vdash \iota.f = v, fh \approx \iota.f(v); v, h} \text{C-WK-WRITE-BODY}$	
$\frac{}{\varepsilon \vdash new\ c, fh \approx new\ c.init_c(null), h} \text{C-WK-NEW}$	
$\frac{fh(\iota) \text{ is undefined} \quad h(\iota) = c}{\iota \vdash new\ c, fh \approx \iota.init_c(null), h} \text{C-WK-NEW-BEGIN}$	
$\frac{h(\iota) = c \quad fh(\iota) \text{ is undefined} \quad fP\downarrow_4(c) = \{f_1\ _, \dots, f_k\ _ \} \quad r < k}{\iota \vdash new\ c, fh \approx \iota.f_r(null); \dots; \iota.f_k(null); \iota, \iota.f_1(null), \dots, \iota.f_{r-1}(null), h} \text{C-WK-NEW-BODY}$	
$\frac{\varphi \vdash f_e, fh \approx e, \bar{e}, h \quad \varphi' \vdash f_{e'}, fh \approx e', \bar{e}', h \quad f_e = v \text{ or } \varphi' = \varepsilon}{\varphi \oplus \varphi' \vdash f_e.m(f_{e'}), fh \approx e.m(e'), \bar{e}, \bar{e}', h} \text{C-WK-METH}$	

Figure 15: Weak Correspondence for Expressions

is being initialised by the *init...* chord the form of φ is ι ; when no field or object is being accessed then φ is empty, denoted ε . Hence, $\varphi \in Addr \times Id^f \cup Addr \cup \{\varepsilon\}$.

In figure 16 we illustrate the use of weak correspondence on the original example from figure 13, where an f SCHOOL expression writes the value 2 to the *permits* field. It is now possible to apply a weak correspondence judgement after each step of the SCHOOL execution, including intermediate configurations where the f SCHOOL configuration has not executed (denoted by a dotted line).

The initial configurations weakly correspond through the C-WK-WRITE rule, whose premises are satisfied through the C-WK-VAL rule, applied both on the receiver address ι_0 and on the value 2, which appear in both configurations.

After the initial SCHOOL configuration evaluates through the JOIN rule (the f SCHOOL configuration does not evaluate), the invocation $\iota_0.permits(3)$ is missing from the configuration (it has been replaced by *voidValue*). Accordingly, φ contains $\iota_0.permits$ and allows us to judge the second SCHOOL configuration

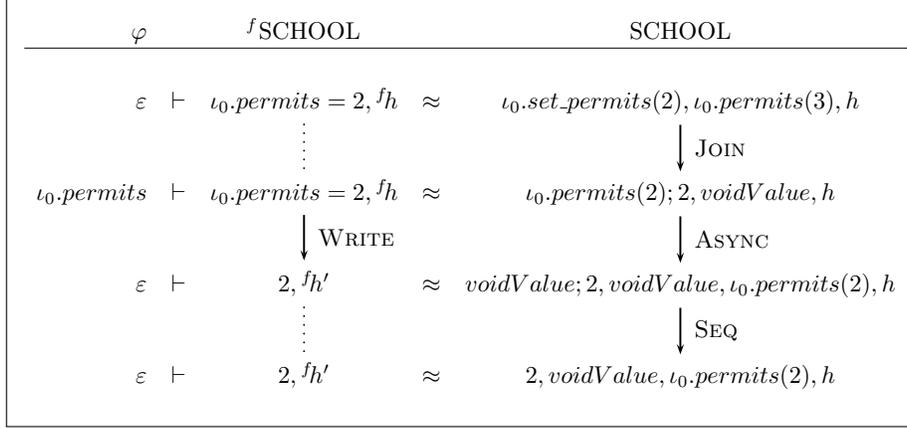


Figure 16: Example of weak correspondence.

as weakly corresponding to the initial f SCHOOL configuration, through the C-WK-WRITE-BODY rule.

The next step of evaluation in SCHOOL is performed via the ASYNC rule, and this time the f SCHOOL configuration makes its single step of evaluation (via the WRITE rule). The invocation representing the new current value of the *permits* field is once again present in the configuration, and accordingly φ becomes empty. Through the C-WK-VOID rule the two new configurations weakly correspond; the premise of the rule is satisfied through the C-WK-VAL rule, applied on the value 2.

The SCHOOL configuration makes a third, and final, step of evaluation (via the SEQ rule), while the f SCHOOL configuration performs no further evaluation; the two final configurations weakly correspond via the C-WK-VAL rule (and hence φ is empty).

When an expression contains sub-expressions (such as the receiver and argument of a method invocation) we combine φ and φ' through the \oplus operation, defined as follows:

$$\varphi \oplus \varphi' = \begin{cases} \varphi & \text{if } \varphi' = \varepsilon \\ \varphi' & \text{if } \varphi = \varepsilon \\ \text{undefined} & \text{otherwise} \end{cases}$$

Notice that the above definition, by requiring that one of the φ s is empty, prohibits sub-expressions from simultaneously accessing fields, even when those fields are distinct.

This restriction on \oplus is desirable because if we allowed sub-expressions to simultaneously access fields we would consider as corresponding SCHOOL configurations which are unreachable. For example, the SCHOOL configuration $E[\iota.f(v); v].m(\iota.f'(v'); v'), h$ is not reachable due to the evaluation order imposed by contexts.

The following weak correspondence rules are between a single f SCHOOL expression and its encoding:

- C-WK-VOID enables us to dispense with the sequential composition of *voidValue* with any other expression.
- C-WK-VAL indicates that ground values immediately correspond, and $\varphi = \varepsilon$ since the SCHOOL expression is either not in a field-access chord body, or is the last term of a field-access chord body and hence has already made the asynchronous invocation representing the field's current value.
- C-WK-READ preserves the weak correspondence for the receiver object of a field-read method call; since the field-read chord has not yet joined the φ remains unchanged.
- C-WK-READ-BODY says that the f SCHOOL term $\iota.f$ corresponds to the sequence of SCHOOL expressions comprising the body of the field-reading chord, and therefore $\varphi = \iota.f$. Since reading a field does not change its value, the rule requires that the value passed as an argument to f , as well as the value returned, must be the value currently mapped to the field in f_h .
- C-WK-WRITE is similar to the C-WK-READ rule; weak correspondence must be preserved for both the receiver and argument, and hence we use the combination $\varphi \oplus \varphi'$. The third premise of this judgement prohibits the receiver and argument from simultaneously accessing fields (as this would allow unreachable states due to the order imposed by contexts) by requiring either the receiver to be a ground value or the argument to not be accessing a field.
- C-WK-WRITE-BODY is similar to the C-WK-READ-BODY rule; however, since the current value of a field is lost during writing, the value mapped to the field in the f SCHOOL heap is ignored.
- C-WK-NEW tells us that the object creation term in f SCHOOL directly corresponds with its encoding, and $\varphi = \varepsilon$ since the object initialisation chord has not joined yet.
- C-WK-NEW-BODY says that the f SCHOOL term for object creation corresponds with a sequence of SCHOOL expressions which consist of the partially evaluated body of the *init...* chord, and the asynchronous invocations to the f methods which have already been evaluated, and spawned through ASYNC. Notice that this is the only case where an f SCHOOL expression corresponds to more than one SCHOOL expression.
- C-WK-METH tells us that a method call in f SCHOOL directly corresponds with its encoding in SCHOOL, as long as the receivers and the arguments also correspond, and hence we use the combination $\varphi \oplus \varphi'$. Similar to C-WK-WRITE above, the receiver and argument are prohibited from simultaneously accessing fields.

$$\boxed{
\begin{array}{c}
\frac{f_h \approx h \quad \varphi_i \vdash f_{e_i}, f_h \approx \bar{e}_i, h \quad \bar{e}f = \mathcal{F}(f_h, \varphi_1, \dots, \varphi_n)}{f_{e_1}, \dots, f_{e_n}, f_h \approx \bar{e}_1, \dots, \bar{e}_n, \bar{e}f, h} \text{C-WK} \\
\frac{\bar{f}_e \cong \bar{f}_{e'} \quad \bar{f}_{e'}, f_h \approx \bar{e}', h \quad \overline{\text{voidValue}} \bar{e}' \cong \bar{e}}{\bar{f}_e, f_h \approx \bar{e}, h} \text{C-WK-PERM} \\
\frac{\forall \iota \in \text{dom}(f_h) : f_h(\iota) = (c, _) \implies h(\iota) = c}{f_h \approx h} \text{C-WK-HEAPS}
\end{array}
}$$

Figure 17: Weak Correspondence for Configurations

5.3.2 Configurations

The form of weak correspondence for configurations (figure 17) is:

$$\bar{f}_e, f_h \approx \bar{e}, f_h$$

and expresses that each of the expressions in \bar{f}_e corresponds with exactly one sequence of expressions in \bar{e} , that no two SCHOOL expressions from \bar{e} are bodies of field-access chords or object initialisation chords on the same field and object, and that all values of fields in f_h are encoded by the presence of asynchronous method invocations in \bar{e} ; the last two are ensured by the following definition.

Definition 9 (Field Encoding)

$$\mathcal{F}(f_h, \varphi_1, \dots, \varphi_n) = \begin{cases} \left\{ \begin{array}{l} \iota.f(v) \mid f_h(\iota) \downarrow_2(f) = v \\ \wedge \forall i \in 1..n : \iota.f \neq \varphi_i \end{array} \right\} & \text{if } i \neq j \implies \text{distinct } \varphi_i, \varphi_j \\ \text{undefined} & \text{otherwise} \end{cases}$$

Two φ s are distinct when they indicate different fields and objects; however, when one φ indicates an entire object then the other φ must not indicate a field of that same object (nor the object itself, of course). Hence, we define distinct φ and φ' as: $\iota.f$ and $\iota'.f'$ are distinct iff $\iota \neq \iota'$ or $f \neq f'$; also ι and $\iota'.f'$ are distinct iff $\iota \neq \iota'$; finally $\iota.f$ and ι' are distinct iff $\iota \neq \iota'$.

Hence the above definition ensures that SCHOOL expressions do not “clash”, in the sense that they are not the bodies for field-access chords for the same object and field, or initialisations for the same object. In order to prove soundness and completeness we use this definition to show that asynchronous methods encoding current values of fields act as mutexes and ensure that *get* and *set* methods cannot access these values simultaneously; specifically, this definition shows us which mutexes are “open” at each step of evaluation by inspecting the currently recorded φ s.

The following weak correspondence rules are between an f SCHOOL and an SCHOOL configuration:

- C-WK imposes the following three requirements:
 1. Each object in $^f h$ exists and has the same class in h , as per the C-WK-HEAPS rule.
 2. Each f SCHOOL expression corresponds to exactly one sequence of SCHOOL expressions, and
 3. For each field f with value v in an object ι in $^f h$, either there exists a corresponding method invocation $\iota.f(v)$ (i.e., $\iota.f \in \varphi$ and \bar{e} contains the appropriate invocation) or one of the SCHOOL expressions is currently executing a related access chord (i.e., $\iota.f = \varphi_i$, for some i , which implies $\varphi_i \vdash ^f e_i, ^f h \approx \bar{e}_i, h$).
- C-WK-PERM extends the correspondence relationship to all possible permutations of the expressions, and allows any number of extra expressions containing *voidValue*.

Consider the example in figure 18 which features two Countdown Latch objects, at addresses ι_0 and ι_1 , with current values 5 and 4, respectively. Weak correspondence allows for the simultaneous joining of two field-access chords on the same field of *two different objects* (and prohibits situations such as that of the example in figure 20, see later section 5.3.4). The first column indicates the values of φ s when application of C-WK rule gives us weak correspondence between the f SCHOOL and SCHOOL configurations in the second and third columns, respectively.

We notice that after two steps of SCHOOL evaluation the φ of the top line is $\iota_0.permits$, while that of the bottom line is $\iota_1.permits$, and we can apply the C-WK judgement. On the other hand, if the f SCHOOL configuration consisted of two concurrent writes to the *permits* fields of the *same* Countdown Latch, we would not be able to join both field-write chords in SCHOOL, as the φ would not be distinct, and hence the configurations would no longer weakly correspond.

5.3.3 Soundness

We show in lemma 4 that encoding preserves weak correspondence, and in theorem 6 we prove the stronger version of theorem 5.

Lemma 4 (Encoding Preserves Weak Correspondence)

- (i) $\forall ^f e, ^f h, h : \varepsilon \vdash ^f e, ^f h \approx \mathbb{C}(^f e), h$
- (ii) $\forall \bar{e}, ^f h, h : ^f h \approx h \implies \bar{e}, ^f h \approx \overline{\mathbb{C}(\bar{e})}, h$

Proof 9 (i) by induction on the structure of $^f e$; (ii) by application of C-WK and use of part (i) and the definition of \mathcal{F} .

Theorem 6 (Soundness of The Encoding with Weak Correspondence)

$$\left. \begin{array}{l} \bar{e}, ^f h \approx \bar{e}, h \\ \bar{e}, ^f h \longrightarrow \bar{e}', ^f h' \end{array} \right\} \implies \exists \bar{e}', h' : \bar{e}, h \longrightarrow^* \bar{e}', h' \wedge \bar{e}', ^f h' \approx \bar{e}', h'$$

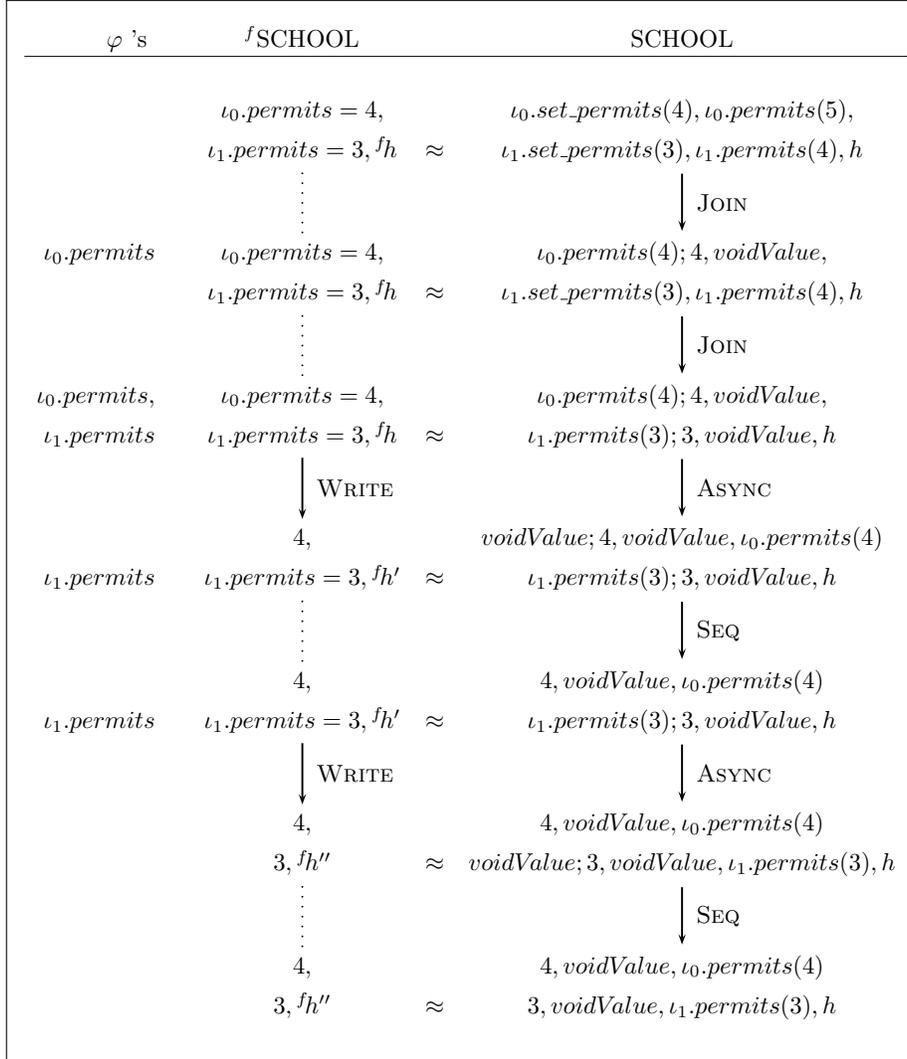


Figure 18: Example of weak correspondence with multiple fields.

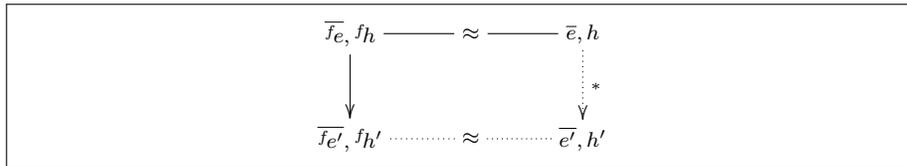


Figure 19: Soundness of the encoding with weak correspondence.

permits of a Countdown Latch at address ι_0 , which currently has the value 5. If its encoding in SCHOOL were to somehow allow both field-write chords to *simultaneously* join, then we would end up with a SCHOOL configuration where *two* invocations of $\iota_0.\text{permits}(_)$ determine the current value of the *permits* field; such a configuration, and all of its further evaluations, would never lead to a strongly corresponding configuration.

When dealing with weak correspondence of an individual f SCHOOL expression we can ignore the corresponding sequence of asynchronous method invocations encoding field values, and hence write $\varphi \vdash ^f e, ^f h \approx e, \bar{e}, h$, where e, \bar{e} is the corresponding sequence of SCHOOL expressions; however, when dealing with the execution of the corresponding SCHOOL configuration we do consider such invocations, and so write $e, \bar{e}, \bar{e}f, h \longrightarrow \bar{e}', h'$, where $\bar{e}f$ is the sequence of these invocations.

Definition 10 below captures the relationship between an f SCHOOL step of evaluation and the current field access, determined by the value of φ . Because in f SCHOOL reading or writing a field is a single step, it is necessary that such a step does not occur *during* the access of a field in the corresponding SCHOOL configuration, as the field-access chord will not yet have replaced the asynchronous invocation encoding the field's current value, and correspondence would break. Hence, such f SCHOOL steps can occur only when φ' is empty.

Definition 10 (f SCHOOL Step)

$$\varphi, ^f e, ^f h \sim \varphi', ^f e', ^f h' \text{ iff } \left\{ \begin{array}{l} ^f e = ^f e' \wedge ^f h = ^f h' \\ \vee \\ ^f e, ^f h \longrightarrow ^f e', ^f h' \quad \wedge \quad \varphi' = \varepsilon \end{array} \right.$$

Using the above definition we show in lemma 5 below that when the SCHOOL configuration weakly corresponding to an individual f SCHOOL expression makes a single step of evaluation, the resulting configuration weakly corresponds to either the original f SCHOOL configuration or to the f SCHOOL configuration resulting from a step of evaluation. Furthermore, the sequence of asynchronous method invocations encoding field current values adheres to definition 9, taking into account when field-access chords return such an invocation and when they join and consume such an invocation.

Lemma 5 (Preservation of Weak Correspondence)

$$\left. \begin{array}{l} \varphi \vdash ^f e, ^f h \approx e, \bar{e}, h \\ \bar{e}f = \mathcal{F}(^f h, \varphi, \varphi_1, \dots, \varphi_k) \\ ^f h \approx h \\ e, \bar{e}, \bar{e}f, h \longrightarrow e', \bar{e}', \bar{e}f', h' \end{array} \right\} \implies \begin{array}{l} \exists \varphi', ^f e', ^f h' : \\ \varphi, ^f e, ^f h \sim \varphi', ^f e', ^f h' \\ \varphi' \vdash ^f e', ^f h' \approx e', \bar{e}', h' \\ \bar{e}f' = \mathcal{F}(^f h', \varphi', \varphi_1, \dots, \varphi_k) \\ ^f h' \approx h' \end{array}$$

Proof 11 *By structural induction on the judgement $\varphi \vdash ^f e, ^f h \approx e, \bar{e}, h$ and use of definition 10; we perform a case analysis on the weak correspondence judgement applied.*

The case for C-WK-VAL is trivial as no evaluation rule is applicable; the case for C-WK-VOID is straightforward through the application of the inductive

hypothesis; the case for C-WK-METH is shown through application of the inductive hypothesis and through the use of the definition for $\varphi_0 \oplus \varphi_1$ (where φ_0 is used for the receiver and φ_1 for the argument), which ensures a valid φ' .

The cases for C-WK-READ and C-WK-WRITE consume an invocation from \overline{ef} , while those for C-WK-READ-BODY and C-WK-WRITE-BODY replace such invocations.

The case for C-WK-NEW is straightforward (care must be taken with heaps). The case for C-WK-NEW-BEGIN has two sub-cases which depend on the number of fields defined in the class of the object being considered: the initialisation chord will result in either a set sequence of initial invocations to field-encoding methods (when one or more fields are defined) or directly result in returning the newly-created object (when no fields are defined); accordingly, the resulting configurations are shown to correspond with C-WK-NEW-BODY in the former case and C-WK-VAL in the latter case.

Finally, the case for C-WK-NEW-BODY also has two sub-cases depending on the number of initial invocations remaining in the execution of the initialisation chord: if a single invocation remains then the current expression becomes the address of the object being considered, and the resulting configurations are shown to correspond through the use of C-WK-VOID and C-WK-VAL; if, however, two or more invocations remain then the same C-WK-NEW-BODY judgement is applied again, with r having increased by one.

We now use weak correspondence to show completeness. Theorem 7 says that for corresponding f SCHOOL and SCHOOL configurations, if the SCHOOL configuration makes one evaluation step, then correspondence is preserved in one of two ways: either the SCHOOL configuration evaluated into one of the intermediate steps and hence still corresponds with the initial f SCHOOL configuration, or the f SCHOOL configuration can make an evaluation step and the two resulting configurations correspond.

Theorem 7 (Completeness of The Encoding)

$$\left. \begin{array}{l} \overline{f_e, f_h} \approx \overline{e, h} \\ \overline{e, h} \longrightarrow \overline{e', h'} \end{array} \right\} \implies \overline{f_e, f_h} \approx \overline{e', h'} \vee \exists \overline{f_{e'}}, \overline{f_{h'}} : \overline{f_e, f_h} \longrightarrow \overline{f_{e'}}, \overline{f_{h'}} \wedge \overline{f_{e'}}, \overline{f_{h'}} \approx \overline{e'}, \overline{h'}$$

Proof 12 *By case analysis on the weak correspondence judgement applied to obtain $\overline{f_e, f_h} \approx \overline{e, h}$, and use of lemma 5 and definition 10. The application of the lemma yields the relationship $\varphi, ^f e, ^f h \sim \varphi', ^f e', ^f h'$ which determines whether the f SCHOOL expression has evaluated or has remained the same; the other three results satisfy the four premises of C-WK (up to permutation and ignoring of leftover voidValue expressions, both of which are handled by C-WK-PERM).*

The left diagram from figure 21 illustrates the above theorem; the solid lines represent the premises, and the dashed and dotted lines represent the first and second conclusions, respectively.

In theorem 5 we had shown soundness for the stronger version of correspondence, \approx , rather than \sim . The converse of theorem 5 does not hold directly, but

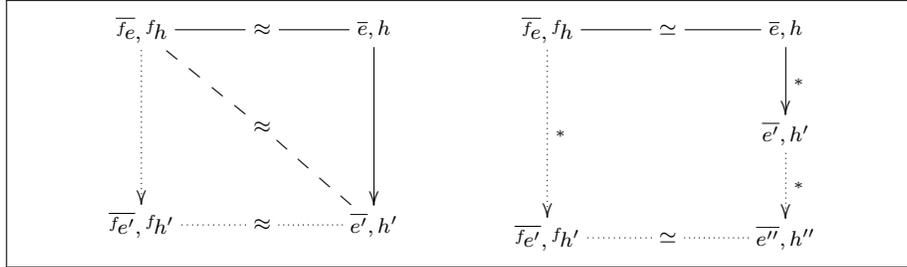


Figure 21: Weak completeness and strong completeness of the encoding.

we can prove the following “diamond” property: if the SCHOOL expressions \bar{e} are the encoding of the f SCHOOL expressions \bar{f}_e , and the expressions \bar{e} evaluate in several steps to some expressions \bar{e}' , then the f SCHOOL expressions can be evaluated to some expressions $\bar{f}_{e'}$ whose encoding is the result of *further* evaluation of \bar{e}' (theorem 8 below).

In order to prove the final theorem it is useful to notice that weak correspondence encompasses strong correspondence, as the latter is a special case of the former where no field-access chord is currently joined (and hence all φ s are empty). In order to see this, consider the collection of weak correspondence judgements for individual expressions which feature an empty φ : this gives correspondence between an expression $^f e$ and its direct encoding $\mathbb{C}(^f e)$ (lemma 6 below), hence matching the strong correspondence judgement C-STRONG for each f SCHOOL expression. Since no field-access chords are currently executing, given by definition 9, the premise of C-STRONG is satisfied. However, when fields *are* being accessed, we must show (lemmas 7 and 8 below) that further evaluation always leads to corresponding configurations where no fields are being accessed any longer, and thus strong correspondence is always reached.

First, we consider weakly-corresponding expressions which do not access fields or objects, and show that the SCHOOL expression is the encoding of its corresponding f SCHOOL expression (lemma 6 below). However, during evaluation it is possible for multiple *voidValue* terms to be present; as these terms will be discarded by the implied sequencing semantics, we define the equality $=_{\text{voidValue}}$ which holds when all *voidValue* terms are removed.

Lemma 6 (Weak Correspondence When No Field Access)

$$\varepsilon \vdash ^f e, ^f h \approx e, h \implies e =_{\text{voidValue}} \mathbb{C}(^f e)$$

Proof 13 *By case analysis on the derivation of $\varepsilon \vdash ^f e, ^f h \approx e, h$.*

Second, we consider weakly-corresponding expressions which access a field or object, and show that for any such collection of expressions from the SCHOOL configuration (hence $\varphi \neq \varepsilon$), it is always the case that they will eventually evaluate so as to cease accessing the field or object, and additionally maintain weak correspondence with the result of evaluating the corresponding f SCHOOL expression (hence $\varphi = \varepsilon$). Furthermore, weak correspondence between the rest

of the expressions in the configurations is preserved, and the resulting heaps correspond (lemma 7 below).

Lemma 7 (Field-Accessing Expressions Complete)

$$\left. \begin{array}{l} \varphi \vdash {}^f e, {}^f h \approx e, \bar{e}, h \\ \forall i \in 1..k : \varphi_i \vdash {}^f e_i, {}^f h \approx e_i, \bar{e}_i, h \\ \bar{e}f = \mathcal{F}({}^f h, \varphi, \varphi_1, \dots, \varphi_k) \\ \varphi \neq \varepsilon \end{array} \right\} \implies$$

$$\begin{array}{l} \exists {}^f e', {}^f h', e', \bar{e}f', h' : \\ {}^f e, {}^f h \longrightarrow {}^f e', {}^f h' \\ e, \bar{e}, \bar{e}f, h \longrightarrow^* e', \bar{e}f', h' \\ \varepsilon \vdash {}^f e', {}^f h' \approx e', h' \\ \forall i \in 1..k : \varphi_i \vdash {}^f e_i, {}^f h' \approx e_i, \bar{e}_i, h' \\ \bar{e}f' = \mathcal{F}({}^f h', \varphi_1, \dots, \varphi_k) \\ {}^f h' \approx h' \end{array}$$

Proof 14 *By structural induction on the weak correspondence judgement $\varphi \vdash {}^f e, {}^f h \approx e, \bar{e}, h$.*

We generalise the above result to weakly corresponding configurations; hence, for any SCHOOL and ^fSCHOOL configurations which weakly correspond, both eventually evaluate so that all fields cease to be accessed and the resulting expressions, as well as the heaps, weakly correspond (lemma 8 below).

Lemma 8 (Field-Accessing Configurations Complete)

$$\bar{e}, {}^f h \approx \bar{e}, h \implies$$

$$\begin{array}{l} \exists \bar{e}', {}^f h', \bar{e}', \bar{e}f', h' : \\ \bar{e}, {}^f h \longrightarrow^* \bar{e}', {}^f h' \\ \bar{e}, h \longrightarrow^* \bar{e}', \bar{e}f', h' \\ \forall i \in 1..n : \varepsilon \vdash {}^f e'_i, {}^f h' \approx e'_i, h' \\ \bar{e}f' = \mathcal{F}({}^f h') \\ {}^f h' \approx h' \end{array}$$

Proof 15 *By induction on the number of expressions in \bar{e} and repeated application of lemma 7.*

Finally, we use theorem 7 and the above lemmas to show strong completeness of the encoding (theorem 8 below).

Theorem 8 (Strong Completeness of The Encoding)

$$\left. \begin{array}{l} \bar{e}, {}^f h \simeq \bar{e}, h \\ \bar{e}, h \longrightarrow^* \bar{e}', h' \end{array} \right\} \implies \begin{array}{l} \exists \bar{e}', {}^f h', \bar{e}'', h'' : \\ \bar{e}', h' \longrightarrow^* \bar{e}'', h'' \wedge \\ \bar{e}, {}^f h \longrightarrow^* \bar{e}', {}^f h' \wedge \\ \bar{e}', {}^f h' \simeq \bar{e}'', h'' \end{array}$$

Proof 16 We observe that $\overline{f_e, f_h} \simeq \overline{e, h}$ implies $\overline{f_e, f_h} \approx \overline{e, h}$, and hence we can repeatedly apply theorem 7 so as to obtain $\overline{e, h} \xrightarrow{*} \overline{e'}, h'$ and $\overline{f_e, f_h} \xrightarrow{*} \overline{f_{e_0}, f_{h_0}}$ and $\overline{f_{e_0}, f_{h_0}} \approx \overline{e'}, h'$. From lemma 8 we can now obtain $\overline{e'}, h' \xrightarrow{*} \overline{e''}, h''$ and $\overline{f_{e_0}, f_{h_0}} \xrightarrow{*} \overline{f_{e'}, f_{h'}}$ and $\overline{f_{e'}, f_{h'}} \approx \overline{e''}, h''$ where all φ s are empty. Finally, from lemma 6 we obtain $\overline{f_{e'}, f_{h'}} \simeq \overline{e''}, h''$.

The right diagram from figure 21 illustrates the above theorem; solid lines represent the premises and dotted lines represent the conclusions.

In this section we have shown an encoding of programs with fields using only chords. This encoding is sound, since all original program behaviours are preserved, and complete, since no new behaviours are introduced. We have shown correspondence between original and encoded program executions; however, due to concurrency and non-deterministic interleaving of expressions, we allowed a “temporary breaking” of correspondence. Such breaks are permissible, as we have shown that encoded programs reach corresponding configurations again.

6 Related work

Chords first appeared in the experimental language Polyphonic C[#] [4], [5] and then in C_ω [3]. These are both extensions to C[#], so also contain explicit thread operations and the more common synchronisation constructs (e.g. monitors). The interaction between the two concurrency paradigms would benefit from further study. There are also various library-based chorded extensions of languages such as Java [21, 20] and Scala [9].

The ideas for chords originate in the Join-Calculus [14, 15], which was implemented in the JoCaml language [12]. The Join-Calculus itself was heavily influenced by ideas from the Chemical Abstract Machine (CHAM) [6], resulting in the Reflexive Chemical Abstract Machine (RCHAM) [13].

The Join Calculus was designed to be directly implementable as a functional ML-style language, while SCHOOL is a generalised model of imperative, object-oriented languages which feature chords. There are several similarities between the Join-Calculus and SCHOOL, and some fundamental differences.

Overall, both models describe the concurrent and asynchronous execution of processes. Both the Join-Calculus and SCHOOL posit the concept of join patterns as the sole means of synchronisation between processes, and all other constructs and data structures are built on top of the joining mechanism. Join patterns in the Join-Calculus correspond with chord headers in SCHOOL.

An important difference between the two is that asynchronous concurrency is the primitive unit of execution in the Join-Calculus, and synchronous execution is achieved via continuation-passing, while in SCHOOL, execution is synchronous (with evaluation contexts), and asynchrony is obtained through asynchronous chords which become strung and spawn new processes.

A straightforward parallel is that between messages in the Join-Calculus, and invocations of methods in SCHOOL. Messages in the Join-Calculus are sent asynchronously and travel to a known destination, a channel name, which

is statically bound in its immediately enclosing lexical scope. In SCHOOL, asynchronous method invocations are represented as expressions executing in parallel with the rest of the processes. The invocations are sent to known objects via the objects' addresses, and method names are statically bound to their class.

In the Join-Calculus, messages are consumed upon arrival, and paired via pattern-matching to join definitions (selection of a join pattern when more than one is possible is non-deterministic). In SCHOOL, invocations are consumed by the joining of chords, whose selection is based on pattern-matching (when more than one chord can be chosen the result is obtained non-deterministically). Hence both models are characterised by locality in terms of pattern-matching.

One significant effect of the synchronous execution of SCHOOL is that a chord can have the same method of return type void participating as both the synchronous part *and* in the asynchronous part of a chord. Consequently, pattern-matching in SCHOOL is not linear, while the Join-Calculus imposes strict linearity in join patterns.

In terms of a memory model, the Join-Calculus features a global namespace, while in SCHOOL there is a common heap between all processes. As new names can be created in the Join-Calculus, so new heap entries can be created in SCHOOL.

A fundamental difference between the Join-Calculus and SCHOOL is that the former is high-level and based on name-passing (indeed, only names can be passed), while SCHOOL has no high-level constructs and only values can be passed. As a consequence of having only names, the Join-Calculus requires all constructs to be encoded via such names, while SCHOOL features objects as first-class constructs (and ^fSCHOOL SCHOOL features fields as well), and consequently SCHOOL has an inheritance mechanism. Work on the Objective Join-Calculus has resulted in encodings of objects and an inheritance mechanism borrowed from ML (which uses pattern-matching on program definitions and *self*). In both models encodings of constructs necessitate naming restrictions (such as private names in the Join-Calculus and method names such as *init* in SCHOOL).

Probably the closest approach to concurrent language design, outside of the process calculus community, can be seen in Concurrent ML (CML) [18], which allows programmers to define new synchronisation and communication abstractions as first-class, thus giving flexibility to tailor concurrency abstractions.

There are many different ways that concurrency is expressed in object-oriented languages. In Simula [10], for example, objects were processes and messages were passed between them. There are several excellent surveys describing the different paradigms [8, 17].

Probably the oldest model similar to chords are the active objects [7] of actor languages [2], that combine objects and processes. Actors use asynchronous message passing and continuations to provide concurrency, and state is held in the messages, as with chords. However, a message sent to an Actor has an explicit method destination, while with chords, it is the *join* semantics and scheduling which ultimately decide which chord consumes a message.

7 Conclusions and Future Work

The model of chords presented in this paper aims to provide a simple and formal basis for investigating the crucial aspects of chorded language design, with a focus on the semantics of *join* in object-oriented languages.

Our model enables us to reason about the interaction of other language constructs with chords (as we have shown with fields in this paper), and argue about the relative expressive power, and hence utility, of including constructs in a chorded language.

SCHOOL does not attempt to model specific chorded languages, rather, it is intended to be an abstract and flexible model of chorded languages in general. There are valid SCHOOL programs which are not valid in Polyphonic C[#], since SCHOOL abstracts from limitations imposed on Polyphonic C[#] programs for avoiding inheritance issues, restrictions not necessary for proving subject reduction and completeness.

In this paper we have investigated the interaction of chords with fields. We have found that chords suffice to fully encode the behaviour of programs using fields, and hence fields to not add expressive power to chorded languages. In order to demonstrate this we have shown correspondence between ^fSCHOOL, an extension of SCHOOL with fields, and an encoding of fields using only chords. Our encoding preserves the behaviours of original programs and does not introduce new behaviours.

As chords are a construct within a conventional object-oriented language which already contains more traditional concurrency constructs, such as monitors and explicit threads, we would like to obtain similar results for these. Our more long-term ambitions are to investigate and formalise the properties of larger, cross-cutting issues of chorded programs such as exception handling, scheduling, and analysis and verification (such as detection of progress, lack of deadlock or livelock, and so on).

Acknowledgements

We would like to thank John Knottenbelt for his solution to the Countdown Latch, Alex Buckley for his work on the original model of SCHOOL₀, Feruccio Damiani for his suggestions on the formalism, Paola Giannini for her insight which lead to the unified treatment of synchronous and asynchronous chords (and a much simpler semantics without explicit queues), and Neil Datta and Anastasia Niarchou for their general comments.

References

- [1] Martín Abadi and Luca Cardelli. *A Theory of Objects*. Springer-Verlag New York Inc., 1998.

- [2] Gul Agha and Carl Hewitt. Concurrent programming using actors. In *Object-oriented concurrent programming*, pages 37–53. MIT Press, 1987.
- [3] Nick Benton, Gavin Bierman, Luca Cardelli, Erik Meijer, Claudio Russo, and Wolfram Schulte. C_ω . <http://research.microsoft.com/Comega/>, 2004.
- [4] Nick Benton, Luca Cardelli, and Cédric Fournet. Modern concurrency abstractions for C^\sharp . In B. Magnusson, editor, *Proc. of ECOOP02*, volume 2374, pages 415–440. Springer Press, 2002.
- [5] Nick Benton, Luca Cardelli, and Cédric Fournet. Modern concurrency abstractions for C^\sharp . *ACM Trans. Program. Lang. Syst.*, 26(5):769–804, 2004.
- [6] Gerard Berry and Gerard Boudol. The chemical abstract machine. In *POPL '90: Proceedings of the 17th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 81–94. ACM Press, 1990.
- [7] Andrew Black, Norman Hutchinson, Eric Jul, and Henry Levy. Object structure in the emerald system. In *OOPLSA '86: Conference proceedings on Object-oriented programming systems, languages and applications*, pages 78–86. ACM Press, 1986.
- [8] Jean-Pierre Briot, Rachid Guerraoui, and Klaus-Peter Lohr. Concurrency and distribution in object-oriented programming. *ACM Comput. Surv.*, 30(3):291–329, 1998.
- [9] Vincent Cremet. Join definitions in scala. <http://lamp.epfl.ch/cremet/join.in.scala/>.
- [10] Ole-Johan Dahl and Kristen Nygaard. Simula: an algol-based simulation language. *Commun. ACM*, 9(9):671–678, 1966.
- [11] Sophia Drossopoulou, Alexis Petrounias, Alex Buckley, and Susan Eisenbach. SCHOOL: a small chorded object-oriented language. In Maribel Fernández and Ian Machie, editors, *Proceedings of ICALP '05 First International Workshop on Developments in Computational Models*, pages 55–64, 2005.
- [12] Cédric Fournet, Fabrice Le Fessant, Luc Maranget, and Alain Schmitt. Jocaml: a language for concurrent distributed and mobile programming. In Johan Jeuring and Simon Peyton Jones, editors, *Advanced Functional Programming: 4th International School, AFP 2002*, volume 2638 of *LNCS*. Springer-Verlag, 2003.
- [13] Cédric Fournet and Georges Gonthier. The reflexive CHAM and the join-calculus. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 372–385, St. Petersburg Beach, Florida, United States, 1996. ACM Press, New York, USA.

- [14] Cédric Fournet and Georges Gonthier. The join calculus: a language for distributed mobile programming. *Applied Semantics Summer School*, pages 1–66, 2008.
- [15] Cédric Fournet, Georges Gonthier, Jean-Jacques Lévy, Luc Maranget, and Didier Rémy. A calculus of mobile agents. In *Proceedings of the 7th International Conference on Concurrency Theory*, volume 1119 of *Lecture Notes In Computer Science*, pages 406 – 421, 1996.
- [16] Demetrios Nicolaou. A virtual machine for a chorded programming language. Master’s thesis, Department of Computing, Imperia College London, 2007.
- [17] Michael Philippsen. A survey on concurrent object-oriented languages. *Concurrency: Practice and Experience*, 12(10):917–980, August 2000.
- [18] J. H. Reppy. Cml: A higher-order concurrent language. In *Proceedings of the ACM SIGPLAN, Conference on Programming Language Design and Implementation*, pages 293–305, 1991.
- [19] Lefteris Volanakis. Implementation of a chorded compiler. Master’s thesis, Department of Computing, Imperial College London, 2006.
- [20] G Stewart von Itzstein and David Kearney. Applications of join java. In *CRPITS ’02: Proceedings of the seventh Asia-Pacific conference on Computer systems architecture*, pages 37–46. Australian Computer Society, Inc., 2002.
- [21] Tim Wood. A chorded compiler for java. Master’s thesis, Imperial College, London, June 2004.