

Verifying Temporal Properties of Reactive Systems by Transformation

G.W. Hamilton

School of Computing and Lero
Dublin City University
Ireland

hamilton@computing.dcu.ie

We show how program transformation techniques can be used for the verification of both safety and liveness properties of reactive systems. In particular, we show how the program transformation technique *distillation* can be used to transform reactive systems specified in a functional language into a simplified form that can subsequently be analysed to verify temporal properties of the systems. Example systems which are intended to model mutual exclusion are analysed using these techniques with respect to both safety (mutual exclusion) and liveness (non-starvation), with the errors they contain being correctly identified.

1 Introduction

Formal verification of software components is gaining more and more prominence as a viable methodology for increasing the reliability and reducing the cost of software production. We consider here the problem of verifying properties of *reactive systems*, i.e., systems which continuously react to external events by changing their internal state and producing outputs. The properties of such systems are usually expressed using a *temporal logic* such as Computational Tree Logic (CTL) or Linear-time Temporal Logic (LTL). These logics are used to express *safety* properties which essentially state that nothing bad will happen, and *liveness* properties which essentially state that something good will eventually happen.

Model checking is a well established technique originally developed for the verification of temporal properties of finite state systems [4]. However, reactive systems usually have an infinite number of states. Model checking techniques therefore need to be extended to handle such systems, but the problem of verifying such systems is undecidable in general. Most proposed approaches to this problem are semi-automatic and involve either mathematical (*co*-)induction [3, 8] or *abstraction* to finite state models [11, 16]. Fold/unfold program transformation techniques have more recently been proposed as an automatic approach to this problem. Folding corresponds to the application of a (co-)inductive hypothesis and generalisation corresponds to abstraction. Many such techniques have been developed for logic programs (e.g. [13, 17, 5, 1, 9]). However, very few such techniques have been developed for functional programs (with the work of Lisitsa and Nemytykh [14, 2] using supercompilation [19] being a notable exception), and these deal only with safety properties.

In this paper, we describe a fold/unfold program transformation technique which can be used to verify both safety and liveness properties of reactive systems which have been specified using a functional language. The program transformation technique which we use is our own *distillation* [6, 7] which builds on top of positive supercompilation [18], but is much more powerful. Distillation is used to transform programs into a simplified form which makes them much easier to analyse. We argue that since distillation removes more intermediate structures than positive supercompilation more accurate results are obtained; these intermediate structures need to be generalised in positive supercompilation,

leading to a loss of information. We define a number of verification rules on the simplified form produced by distillation to verify temporal formulae. These techniques are then applied to a number of example systems which are intended to model mutually exclusive access to a critical resource by two processes, revealing a number of errors.

The remainder of this paper is structured as follows. In Section 2, we introduce the functional language over which our verification techniques are defined. In Section 3, we show how to model reactive systems in our language, and give a number of example systems which are intended to model mutually exclusive access to a critical resource by two processes. In Section 4, we describe how we specify temporal properties, and specify both safety (mutual exclusion) and liveness (non-starvation) for the example systems. In Section 5, we describe our technique for verifying temporal properties of reactive systems and apply this technique to the example systems to verify the previously specified temporal properties. Section 6 concludes and considers related work.

2 Language

In this section, we describe the syntax and semantics of the higher-order functional language which will be used throughout this paper.

2.1 Syntax

The syntax of our language is given in Figure 1.

$e ::= x$	Variable
$c e_1 \dots e_k$	Constructor Application
$\lambda x. e$	λ -Abstraction
f	Function Call
$e_0 e_1$	Application
case e_0 of $p_1 \rightarrow e_1 \mid \dots \mid p_k \rightarrow e_k$	Case Expression
let $x = e_0$ in e_1	Let Expression
e_0 where $f_1 = e_1 \dots f_n = e_n$	Local Function Definitions
$p ::= c x_1 \dots x_k$	Pattern

Figure 1: Language Grammar

A program in the language is an expression which can be a variable, constructor application, λ -abstraction, function call, application, **case**, **let** or **where**. Variables introduced by λ -abstractions, **let** expressions and **case** patterns are *bound*; all other variables are *free*. An expression which contains no free variables is said to be *closed*.

Each constructor has a fixed arity; for example *Nil* has arity 0 and *Cons* has arity 2. In an expression $c e_1 \dots e_n$, n must equal the arity of c . The patterns in **case** expressions may not be nested. No variable may appear more than once within a pattern. We assume that the patterns in a **case** expression are non-overlapping and exhaustive. We also allow a wildcard pattern $_$ which always matches if none of the earlier patterns match. Types are defined using algebraic data types, and it is assumed that programs are well-typed. Erroneous terms such as **case** $(\lambda x. e)$ **of** $p_1 \rightarrow e_1 \mid \dots \mid p_k \rightarrow e_k$ and $(c e_1 \dots e_n) e$ where c is of arity n cannot therefore occur.

2.2 Semantics

The call-by-name operational semantics of our language is standard: we define an evaluation relation \Downarrow between closed expressions and *values*, where values are expressions in *weak head normal form* (i.e. constructor applications or λ -abstractions). We define a one-step reduction relation $\overset{r}{\rightsquigarrow}$ inductively as shown in Figure 2, where the reduction r can be f (unfolding of function f), c (elimination of constructor c) or β (β -substitution).

$$\begin{array}{c}
((\lambda x.e_0) e_1) \overset{\beta}{\rightsquigarrow} (e_0\{x \mapsto e_1\}) \quad (\mathbf{let} \ x = e_0 \ \mathbf{in} \ e_1) \overset{\beta}{\rightsquigarrow} (e_1\{x \mapsto e_0\}) \\
\\
\frac{f = e}{f \overset{f}{\rightsquigarrow} e} \qquad \frac{e_0 \overset{r}{\rightsquigarrow} e'_0}{(e_0 e_1) \overset{r}{\rightsquigarrow} (e'_0 e_1)} \\
\\
\frac{p_i = c \ x_1 \dots x_n}{(\mathbf{case} \ (c \ e_1 \dots e_n) \ \mathbf{of} \ p_1 : e'_1 \mid \dots \mid p_k : e'_k) \overset{c}{\rightsquigarrow} (e_i\{x_1 \mapsto e_1, \dots, x_n \mapsto e_n\})} \\
\\
\frac{e_0 \overset{r}{\rightsquigarrow} e'_0}{(\mathbf{case} \ e_0 \ \mathbf{of} \ p_1 : e_1 \mid \dots \mid p_k : e_k) \overset{r}{\rightsquigarrow} (\mathbf{case} \ e'_0 \ \mathbf{of} \ p_1 : e_1 \mid \dots \mid p_k : e_k)}
\end{array}$$

Figure 2: One-Step Reduction Relation

We use the notation $e \overset{r}{\rightsquigarrow}$ if the expression e reduces, $e \Uparrow$ if e diverges, $e \Downarrow$ if e converges and $e \Downarrow v$ if e evaluates to the value v . These are defined as follows, where $\overset{r}{\rightsquigarrow}^*$ denotes the reflexive transitive closure of $\overset{r}{\rightsquigarrow}$:

$$\begin{array}{ll}
e \overset{r}{\rightsquigarrow}, \text{ iff } \exists e'. e \overset{r}{\rightsquigarrow} e' & e \Downarrow, \text{ iff } \exists v. e \Downarrow v \\
e \Downarrow v, \text{ iff } e \overset{r}{\rightsquigarrow}^* v \wedge \neg(v \overset{r}{\rightsquigarrow}) & e \Uparrow, \text{ iff } \forall e'. e \overset{r}{\rightsquigarrow}^* e' \Rightarrow e' \overset{r}{\rightsquigarrow}
\end{array}$$

3 Specifying Reactive Systems

In this section, we show how to specify reactive systems in our programming language. While reactive systems are usually specified using *labelled transitions systems*, our specifications can be trivially derived from these. Reactive systems have to react to a series of *external events* by updating their *state*. In order to facilitate this, we make use of a *stream* datatype, which is defined as follows:

$$\mathit{Stream} \ a ::= \mathit{Cons} \ a \ \mathit{Stream}$$

A stream is therefore an infinite list of elements of type a . Our programs will map an input stream of external events and an initial state to an output stream of states. In this paper, we wish to analyse a number of systems which are intended to implement mutually exclusive access to a critical resource for two processes. In all of these systems, the external events belong to the following datatype:

$$\mathit{Event} ::= \mathit{Request}_1 \mid \mathit{Request}_2 \mid \mathit{Take}_1 \mid \mathit{Take}_2 \mid \mathit{Release}_1 \mid \mathit{Release}_2$$

Each of the two processes can therefore request access to the critical resource, and take and release this resource. States in all of our example systems belong to the following datatype:

$$\mathit{SysState} ::= \mathit{State} \ \mathit{ProcState} \ \mathit{ProcState}$$

$$ProcState ::= T \mid W \mid U$$

Each process can therefore be thinking (T), waiting for the critical resource (W) or using the critical resource (U). In all of the following examples, the variable es represents the external event stream, and s_1 and s_2 represent the states of the two processes respectively.

Example 1 In the first example shown in Figure 3, each process can request access to the critical resource if neither process is using it, take the critical resource if it is waiting for it, and release the critical resource if it is using it.

```

f es T T
where
f = λes s1 s2.Cons (SysState s1 s2) (case es of
    Cons e es → case e of
        Request1 → case s1 of
            U → f es s1 s2
            | _ → case s2 of
                U → f es s1 s2
                | _ → f es W s2
        | Request2 → case s2 of
            U → f es s1 s2
            | _ → case s1 of
                U → f es s1 s2
                | _ → f es s1 W
        | Take1 → case s1 of
            W → f es U s2
            | _ → f es s1 s2
        | Take2 → case s2 of
            W → f es s1 U
            | _ → f es s1 s2
        | Release1 → case s1 of
            U → f es T s2
            | _ → f es s1 s2
        | Release2 → case s2 of
            U → f es s1 T
            | _ → f es s1 s2)

```

Figure 3: Example 1

Example 2 In the second example shown in Figure 4, each process can request access to the critical resource if it is thinking, take the critical resource if it is waiting for it and the other process is thinking, and release the critical resource if it is using it.

Example 3 In the final example in Figure 5, we implement Lamport's bakery algorithm [10] for two processes. In this example, to request access to the critical resource, each process must take a 'ticket' with a number, and the process with the lowest valued ticket is given precedence. A ticket value of zero indicates that a process has not requested to use the critical resource, so when a process releases the

$f \text{ es } T T$

where

$f = \lambda \text{ es } s_1 s_2. \text{Cons } (\text{SysState } s_1 s_2) (\text{case es of}$
 $\text{Cons } e \text{ es} \rightarrow \text{case } e \text{ of}$
 $\text{Request}_1 \rightarrow \text{case } s_1 \text{ of}$
 $\quad T \rightarrow f \text{ es } W s_2$
 $\quad | _ \rightarrow f \text{ es } s_1 s_2$
 $| \text{Request}_2 \rightarrow \text{case } s_2 \text{ of}$
 $\quad T \rightarrow f \text{ es } s_1 W$
 $\quad | _ \rightarrow f \text{ es } s_1 s_2$
 $| \text{Take}_1 \rightarrow \text{case } s_1 \text{ of}$
 $\quad W \rightarrow \text{case } s_2 \text{ of}$
 $\quad\quad T \rightarrow f \text{ es } U s_2$
 $\quad\quad | _ \rightarrow f \text{ es } s_1 s_2$
 $\quad | _ \rightarrow f \text{ es } s_1 s_2$
 $| \text{Take}_2 \rightarrow \text{case } s_2 \text{ of}$
 $\quad W \rightarrow \text{case } s_1 \text{ of}$
 $\quad\quad T \rightarrow f \text{ es } s_1 U$
 $\quad\quad | _ \rightarrow f \text{ es } s_1 s_2$
 $\quad | _ \rightarrow f \text{ es } s_1 s_2$
 $| \text{Release}_1 \rightarrow \text{case } s_1 \text{ of}$
 $\quad U \rightarrow f \text{ es } T s_2$
 $\quad | _ \rightarrow f \text{ es } s_1 s_2$
 $| \text{Release}_2 \rightarrow \text{case } s_2 \text{ of}$
 $\quad U \rightarrow f \text{ es } s_1 T$
 $\quad | _ \rightarrow f \text{ es } s_1 s_2)$

Figure 4: Example 2

critical resource its ticket value is reset to zero. We therefore add two further variables t_1 and t_2 which give the current ticket number for each process. These are natural numbers belonging to the following datatype:

$$\text{Nat} ::= \text{Zero} \mid \text{Succ Nat}$$

Note that, since there is no limit to the number of a ticket, this is an example of an infinite state system which can cause problems for some model checkers.

4 Specification of Temporal Properties

In this section, we describe how temporal properties of reactive systems defined in our functional language are specified. We use Linear-time Temporal Logic (LTL), in which the set of well-founded formulae (WFF) are defined inductively as follows. All atomic propositions p are in WFF; if φ and ψ are in WFF, then so are:

- $\neg\varphi$
- $\varphi \vee \psi$

$f\ es\ T\ T\ Zero\ Zero$

where

$f = \lambda es\ s_1\ s_2\ t_1\ t_2. Cons\ (SysState\ s_1\ s_2)$

(case es of

$Cons\ e\ es \rightarrow$ **case e of**

$Request_1 \rightarrow$ **case s₁ of**

$T \rightarrow f\ es\ W\ s_2\ (Succ\ t_2)\ t_2$

$| _ \rightarrow f\ es\ s_1\ s_2\ t_1\ t_2$

$| Request_2 \rightarrow$ **case s₂ of**

$T \rightarrow f\ es\ s_1\ W\ t_1\ (Succ\ t_1)$

$| _ \rightarrow f\ es\ s_1\ s_2\ t_1\ t_2$

$| Take_1 \rightarrow$ **case s₁ of**

$W \rightarrow$ **case s₂ of**

$T \rightarrow f\ es\ U\ s_2\ t_1\ t_2$

$| _ \rightarrow$ **case (t₁ < t₂) of**

$True \rightarrow f\ es\ U\ s_2\ t_1\ t_2$

$| False \rightarrow f\ es\ s_1\ s_2\ t_1\ t_2$

$| _ \rightarrow f\ es\ s_1\ s_2\ t_1\ t_2$

$| Take_2 \rightarrow$ **case s₂ of**

$W \rightarrow$ **case s₁ of**

$T \rightarrow f\ es\ s_1\ U\ t_1\ t_2$

$| _ \rightarrow$ **case (t₂ < t₁) of**

$True \rightarrow f\ es\ s_1\ U\ t_1\ t_2$

$| False \rightarrow f\ es\ s_1\ s_2\ t_1\ t_2$

$| _ \rightarrow f\ es\ s_1\ s_2\ t_1\ t_2$

$| Release_1 \rightarrow$ **case s₁ of**

$U \rightarrow f\ es\ T\ s_2\ Zero\ t_2$

$| _ \rightarrow f\ es\ s_1\ s_2\ t_1\ t_2$

$| Release_2 \rightarrow$ **case s₂ of**

$U \rightarrow f\ es\ s_1\ T\ t_1\ Zero$

$| _ \rightarrow f\ es\ s_1\ s_2\ t_1\ t_2$)

Figure 5: Example 3

- $\varphi \wedge \psi$
- $\varphi \Rightarrow \psi$
- $\Box\varphi$
- $\Diamond\varphi$
- $\bigcirc\varphi$

The temporal operator $\Box\varphi$ means that φ is *always* true; this is used to express *safety* properties. The temporal operator $\Diamond\varphi$ means that φ will *eventually* be true; this is used to express *liveness* properties. The temporal operator $\bigcirc\varphi$ means that φ is true in the *next* state. These modalities can be combined to obtain new modalities; for example, $\Box\Diamond\varphi$ means that φ is true infinitely often, and $\Diamond\Box\varphi$ means that φ is eventually true forever. Fairness constraints can also be specified for some external events (those

belonging to the set F) which require that they occur infinitely often. For the examples given in this paper, it is assumed that all external events belong to F .

Propositional models for linear-time temporal formulas consist of an infinite sequence of states $\pi = \langle s_0, s_1, \dots \rangle$ such that each state s_i supplies an assignment to the atomic propositions. The satisfaction relation is extended to formulas in LTL for a model π and position i as follows.

$$\begin{array}{ll}
\pi, i \models p & \text{iff } p \in s_i \\
\pi, i \models \neg\phi & \text{iff } \pi, i \not\models \phi \\
\pi, i \models \phi \vee \psi & \text{iff } \pi, i \models \phi \text{ or } \pi, i \models \psi \\
\pi, i \models \phi \wedge \psi & \text{iff } \pi, i \models \phi \text{ and } \pi, i \models \psi \\
\pi, i \models \phi \Rightarrow \psi & \text{iff } \pi, i \not\models \phi \text{ or } \pi, i \models \psi \\
\pi, i \models \Box\phi & \text{iff } \forall j \geq i. \pi, j \models \phi \\
\pi, i \models \Diamond\phi & \text{iff } \exists j \geq i. \pi, j \models \phi \\
\pi, i \models \bigcirc\phi & \text{iff } \pi, i+1 \models \phi
\end{array}$$

A formula ϕ holds in model π if it holds at position 0 i.e. $\pi, 0 \models \phi$.

The atomic propositions of these temporal formulae can be trivially translated into our functional language. For our verification rules, we define the following datatype for truth values:

$$TruthVal ::= True \mid False \mid Undefined$$

The reason that we use a three-valued logic is that our verification rules must always return an answer, but some of the properties to be verified may be undecidable. For our example programs which attempt to implement mutual exclusion, the following two properties are defined. Within these temporal properties, we use the variable s to denote the current state whose properties are being specified.

Property 1 (Mutual Exclusion) This is a safety property which specifies that both processes cannot be using the critical resource at the same time. This can be specified as follows:

$$\begin{array}{l}
\Box(\text{case } s \text{ of} \\
\quad SysState \ s_1 \ s_2 \ \rightarrow \ \text{case } s_1 \ \text{of} \\
\qquad U \ \rightarrow \ \text{case } s_2 \ \text{of} \\
\qquad \quad U \ \rightarrow \ False \\
\qquad \quad | _ \ \rightarrow \ True \\
\qquad | _ \ \rightarrow \ True)
\end{array}$$

Property 2 (Non-Starvation) This is a liveness property which specifies that each process must eventually get to use the critical resource if they are waiting for it. This can be specified for process 1 as follows:

$$\begin{array}{l}
\Box((\text{case } s \ \text{of} \\
\quad SysState \ s_1 \ s_2 \ \rightarrow \ \text{case } s_1 \ \text{of} \\
\qquad W \ \rightarrow \ True \\
\qquad | _ \ \rightarrow \ False) \Rightarrow \Diamond(\text{case } s \ \text{of} \\
\qquad \qquad SysState \ s_1 \ s_2 \ \rightarrow \ \text{case } s_1 \ \text{of} \\
\qquad \qquad \quad U \ \rightarrow \ True \\
\qquad \qquad \quad | _ \ \rightarrow \ False))
\end{array}$$

The specification of this property for process 2 is similar.

5 Verification of Temporal Properties

In this section, we show how temporal properties of reactive systems defined in our functional language can be verified. To facilitate this, we first of all transform the reactive systems definitions into a simplified form using distillation [6, 7], a powerful program transformation technique which builds on top of the supercompilation transformation [19, 18]. Due to the nature of the reactive systems definitions, in which the input is an external event stream, and the output is a stream of states, the programs resulting from this transformation will take the form e^θ , where e^ρ is defined as follows.

$$\begin{aligned}
 e^\rho & ::= \text{Cons } e_0^\rho \ e_1^\rho \\
 & \quad | \ f \ x_1 \dots x_n \\
 & \quad | \ \mathbf{case} \ x \ \mathbf{of} \ p_1 \rightarrow e_1^\rho \ | \dots \ | \ p_k \rightarrow e_n^\rho, \text{ where } x \notin \rho \\
 & \quad | \ x \ e_1^\rho \dots e_n^\rho, \text{ where } x \in \rho \\
 & \quad | \ \mathbf{let} \ x = \lambda x_1 \dots x_n. e_0^\rho \ \mathbf{in} \ e_1^{\rho \cup \{x\}} \\
 & \quad | \ e_0^\rho \ \mathbf{where} \ f_1 = \lambda x_{1_1} \dots x_{1_k}. e_1^\rho \dots f_n = \lambda x_{n_1} \dots x_{n_k}. e_n^\rho
 \end{aligned}$$

The **let** expressions indicate where generalisation has taken place. The **let** variables are added to the set ρ , and will not be used in the selectors of **case** expressions. These **let** variables will be abstracted so no information can be assumed about them during verification.

We define our verification rules on this restricted form of program as shown in Figure 6. The parameter φ denotes the property to be verified and ϕ denotes the function variable environment. ρ denotes the set of function calls previously encountered; this is used for the detection of loops to ensure termination. ρ is also used in the verification of the \square operator (which evaluates to *True* on encountering a loop), and the verification of the \diamond operator (which evaluates to *False* on encountering a loop); ρ is reset to empty when the verification moves inside these temporal operators. For all other temporal formulae, the value *Undefined* is returned on encountering a loop.

The verification rules can be explained as follows. The logical connectives \wedge , \vee , \Rightarrow and \neg are defined in the usual way for a three-valued logic in our language in rules (1-4). Rules (5a-d) deal with a constructed stream of states. In rule (5a), if we are trying to verify that a property is always true, then we verify that it is true for the first state (with ρ reset to empty) and is always true in all remaining states. In rule (5b), if we are trying to verify that a property is eventually true, then we verify that it is either true for the first state (with ρ reset to empty) or is eventually true in all remaining states. In rule (5c), if we are trying to verify that a property is true in the next state then we verify that the property is true for the next state. In rule (5d), if we are trying to verify that a property is true in the current state then we verify that the property is true for the current state by evaluating the property using the value of the current state for the state variable s . Rules (6a-c) deal with function calls. In rule (6a), if we are trying to verify that a property is always true, then if the function call has been encountered before while trying to verify the same property we can return the value *True*; this corresponds to the standard greatest fixed point calculation normally used for the \square operator in which the property is initially assumed to be *True* for all states. Otherwise, the function is unfolded and added to the set of previously encountered function calls for this property. In rule (6b), if we are trying to verify that a property is eventually true, then if the function call has been encountered before while trying to verify the same property we can return the value *False*; this corresponds to the standard least fixed point calculation normally used for the \diamond property in which the property is initially assumed to be *False* for all states. Otherwise, the function is unfolded and added to the set of previously encountered function calls for this property. In rule (6c), if we are trying to verify that any other property is true, then if the function call has been encountered before we can return the value *Undefined* since a loop has been detected. Otherwise, the function is unfolded and added to the

$$\begin{aligned}
(1) \quad \mathcal{P}[[e]] (\varphi \wedge \psi) \phi \rho &= \mathbf{case} (\mathcal{P}[[e]] \varphi \phi \rho) \mathbf{of} \\
&\quad \mathit{True} \quad \rightarrow \mathcal{P}[[e]] \psi \phi \rho \\
&\quad | \mathit{False} \quad \rightarrow \mathit{False} \\
&\quad | \mathit{Undefined} \rightarrow \mathit{Undefined} \\
(2) \quad \mathcal{P}[[e]] (\varphi \vee \psi) \phi \rho &= \mathbf{case} (\mathcal{P}[[e]] \varphi \phi \rho) \mathbf{of} \\
&\quad \mathit{True} \quad \rightarrow \mathit{True} \\
&\quad | \mathit{False} \quad \rightarrow \mathcal{P}[[e]] \psi \phi \rho \\
&\quad | \mathit{Undefined} \rightarrow \mathcal{P}[[e]] \psi \phi \rho \\
(3) \quad \mathcal{P}[[e]] (\varphi \Rightarrow \psi) \phi \rho &= \mathbf{case} (\mathcal{P}[[e]] \varphi \phi \rho) \mathbf{of} \\
&\quad \mathit{True} \quad \rightarrow \mathcal{P}[[e]] \psi \phi \rho \\
&\quad | \mathit{False} \quad \rightarrow \mathit{True} \\
&\quad | \mathit{Undefined} \rightarrow \mathit{Undefined} \\
(4) \quad \mathcal{P}[[e]] (\neg \varphi) \phi \rho &= \mathbf{case} (\mathcal{P}[[e]] \varphi \phi \rho) \mathbf{of} \\
&\quad \mathit{True} \quad \rightarrow \mathit{False} \\
&\quad | \mathit{False} \quad \rightarrow \mathit{True} \\
&\quad | \mathit{Undefined} \rightarrow \mathit{Undefined} \\
(5a) \quad \mathcal{P}[[\mathit{Cons} \ e_0 \ e_I]] (\Box \varphi) \phi \rho &= \mathcal{P}[[\mathit{Cons} \ e_0 \ e_I]] \varphi \phi \emptyset \wedge \mathcal{P}[[e_I]] (\Box \varphi) \phi \rho \\
(5b) \quad \mathcal{P}[[\mathit{Cons} \ e_0 \ e_I]] (\Diamond \varphi) \phi \rho &= \mathcal{P}[[\mathit{Cons} \ e_0 \ e_I]] \varphi \phi \emptyset \vee \mathcal{P}[[e_I]] (\Diamond \varphi) \phi \rho \\
(5c) \quad \mathcal{P}[[\mathit{Cons} \ e_0 \ e_I]] (\bigcirc \varphi) \phi \rho &= \mathcal{P}[[e_I]] \varphi \phi \rho \\
(5d) \quad \mathcal{P}[[\mathit{Cons} \ e_0 \ e_I]] \varphi \phi \rho &= v, \text{ where } \varphi[e_0/s] \Downarrow v \\
(6a) \quad \mathcal{P}[[f \ x_1 \dots x_n]] (\Box \varphi) \phi \rho &= \begin{cases} \mathit{True}, & \text{if } f \in \rho \\ \mathcal{P}[[e[x_1/x'_1, \dots, x_n/x'_n]]] (\Box \varphi) \phi (\rho \cup \{f\}), & \text{otherwise} \end{cases} \\
&\quad \text{where } \phi(f) = \lambda x'_1 \dots x'_n. e \\
(6b) \quad \mathcal{P}[[f \ x_1 \dots x_n]] (\Diamond \varphi) \phi \rho &= \begin{cases} \mathit{False}, & \text{if } f \in \rho \\ \mathcal{P}[[e[x_1/x'_1, \dots, x_n/x'_n]]] (\Diamond \varphi) \phi (\rho \cup \{f\}), & \text{otherwise} \end{cases} \\
&\quad \text{where } \phi(f) = \lambda x'_1 \dots x'_n. e \\
(6c) \quad \mathcal{P}[[f \ x_1 \dots x_n]] \varphi \phi \rho &= \begin{cases} \mathit{Undefined}, & \text{if } f \in \rho \\ \mathcal{P}[[e[x_1/x'_1, \dots, x_n/x'_n]]] \varphi \phi (\rho \cup \{f\}), & \text{otherwise} \end{cases} \\
&\quad \text{where } \phi(f) = \lambda x'_1 \dots x'_n. e \\
(7a) \quad \mathcal{P}[[\mathbf{case} \ x \ \mathbf{of} \ p_1 \rightarrow e_1 \mid \dots \mid p_n \rightarrow e_n]] (\Diamond \varphi) \phi \rho &= (\bigvee_{p_i \in F} \mathcal{P}[[e_i]] (\Diamond \varphi) \phi \rho) \vee (\bigwedge_{i=1}^n \mathcal{P}[[e_i]] (\Diamond \varphi) \phi \rho) \\
(7b) \quad \mathcal{P}[[\mathbf{case} \ x \ \mathbf{of} \ p_1 \rightarrow e_1 \mid \dots \mid p_n \rightarrow e_n]] \varphi \phi \rho &= \bigwedge_{i=1}^n \mathcal{P}[[e_i]] \varphi \phi \rho \\
(8) \quad \mathcal{P}[[x \ e_1 \dots e_n]] \varphi \phi \rho &= \mathit{Undefined} \\
(9) \quad \mathcal{P}[[\mathbf{let} \ x = e_0 \ \mathbf{in} \ e_I]] \varphi \phi \rho &= \mathcal{P}[[e_I]] \varphi \phi \rho \\
(10) \quad \mathcal{P}[[e_0 \ \mathbf{where} \ f_1 = e_1 \dots f_n = e_n]] \varphi \phi \rho &= \mathcal{P}[[e_0]] \varphi (\phi \cup \{f_1 \mapsto e_1, \dots, f_n \mapsto e_n\}) \rho
\end{aligned}$$

Figure 6: Verification Rules

set of previously encountered function calls. Rules (7a-b) deal with **case** expressions. In rule (7a), if we are trying to verify that a property is eventually true, then we verify that it is either eventually true for at least one of the branches for which there is a fairness assumption (since these branches must be selected

eventually), or that it is eventually true for all branches. In Rule (7b), if we are trying to verify that any other property is true, then we verify that it is true for all branches. In rule (8), if we encounter a free variable, then we return the value *Undefined* since we cannot determine the value of the variable; this must be a **let** variable which has been abstracted, so no information can be determined for it. In rule (9), in order to verify that a property is true for a **let** expression, we verify that it is true for the **let** body; this is where we perform abstraction of the **let** variable. In rule (10), for a **where** expression, the function definitions are added to the environment ϕ .

Theorem 5.1 (Soundness) $\mathcal{P}[[e]] \varphi \ \emptyset \ \emptyset = True \Rightarrow \pi, 0 \models \varphi$, where π is a model for e .

The proof of this is by recursion induction on the verification rules \mathcal{P} .

Theorem 5.2 (Termination) $\forall e \in \text{Prog}, \varphi \in \text{WFF}, \mathcal{P}[[e]] \varphi \ \emptyset \ \emptyset$ always terminates.

Proof of termination is quite straightforward since there will be a finite number of functions and uses of the temporal operators \square and \diamond , and verification of each of these temporal operators will terminate when a function is re-encountered.

Using these rules, we try to verify the two properties (mutual exclusion and non-starvation) for the example programs for mutual exclusion given in Section 3. Firstly, distillation is applied to each of the programs.

Example 1 The result of distilling Example 1 is shown in Figure 7. Verification of Property 1 (mutual exclusion) fails for this transformed program; if the input event stream starts with $[Request_1, Request_2, Take_1, Take_2, \dots]$, then we can see that we end up within the function f_9 , where both processes are using the critical resource.

Example 2 The result of distilling Example 2 is shown in Figure 8. Verification of Property 1 (mutual exclusion) succeeds for this transformed program; we can easily see that there is no state in which both processes are using the critical resource. When trying to prove this always property, as soon as we re-encounter any of the functions within the program, the value *True* is returned by the verification rules. However, verification of Property 2 (non-starvation) fails; if the input event stream starts with $[Request_1, Request_2, \dots]$, then we can see that we end up within the function f_5 . At this point, both processes are waiting for the critical resource, so we need to prove that they will eventually get to use it. When trying to prove this eventuality property, we immediately re-encounter the function f_5 , so the value *False* is returned by the verification rules.

Example 3 The result of distilling Example 3 is shown in Figure 9. We can see that the use of tickets is completely transformed away and that the resulting program has a finite number of states. This is where distillation provides an advantage over other transformation techniques such as positive supercompilation which are not able to remove as many intermediate data structures and thus to transform away the use of tickets. Verification of both Property 1 (mutual exclusion) and Property 2 (non-starvation) succeed for this transformed program. The proof of Property 1 is quite straightforward and similar to the proof of this property for Example 2. If we consider the proof of Property 2 for process 1, if process 1 requests access to the critical resource first then we end up within function f_2 . From this point, process 1 must either take the critical resource immediately moving into function f_4 , or process 2 also requests access to the critical resource moving into function f_6 , at which point process 1 must take the critical resource moving into function f_8 . If process 2 requests access to the critical resource before process 1 then we end up within function f_7 . From this point, process 2 must take the critical resource first, moving into function f_9 , and then eventually release the critical resource moving into function f_2 . From this point, process 1 must eventually take the critical resource, as already shown.

$f_1 es$
where
 $f_1 = \lambda es. Cons (SysState T T) (\text{case } es \text{ of}$
 $Cons e es \rightarrow \text{case } e \text{ of}$
 $Request_1 \rightarrow f_2 es$
 $| Request_2 \rightarrow f_3 es$
 $| _ \rightarrow f_1 es)$
 $f_2 = \lambda es. Cons (SysState W T) (\text{case } es \text{ of}$
 $Cons e es \rightarrow \text{case } e \text{ of}$
 $Take_1 \rightarrow f_4 es$
 $| Request_2 \rightarrow f_5 es$
 $| _ \rightarrow f_2 es)$
 $f_3 = \lambda es. Cons (SysState T W) (\text{case } es \text{ of}$
 $Cons e es \rightarrow \text{case } e \text{ of}$
 $Request_1 \rightarrow f_5 es$
 $| Take_2 \rightarrow f_6 es$
 $| _ \rightarrow f_3 es)$
 $f_4 = \lambda es. Cons (SysState U T) (\text{case } es \text{ of}$
 $Cons e es \rightarrow \text{case } e \text{ of}$
 $Release_1 \rightarrow f_1 es$
 $| _ \rightarrow f_4 es)$
 $f_5 = \lambda es. Cons (SysState W W) (\text{case } es \text{ of}$
 $Cons e es \rightarrow \text{case } e \text{ of}$
 $Take_1 \rightarrow f_7 es$
 $| Take_2 \rightarrow f_8 es$
 $| _ \rightarrow f_5 es)$
 $f_6 = \lambda es. Cons (SysState T U) (\text{case } es \text{ of}$
 $Cons e es \rightarrow \text{case } e \text{ of}$
 $Release_2 \rightarrow f_1 es$
 $| _ \rightarrow f_6 es)$
 $f_7 = \lambda es. Cons (SysState U W) (\text{case } es \text{ of}$
 $Cons e es \rightarrow \text{case } e \text{ of}$
 $Release_1 \rightarrow f_3 es$
 $| Take_2 \rightarrow f_9 es$
 $| _ \rightarrow f_7 es)$
 $f_8 = \lambda es. Cons (SysState W U) (\text{case } es \text{ of}$
 $Cons e es \rightarrow \text{case } e \text{ of}$
 $Release_2 \rightarrow f_2 es$
 $| Take_1 \rightarrow f_9 es$
 $| _ \rightarrow f_8 es)$
 $f_9 = \lambda es. Cons (SysState U U) (\text{case } es \text{ of}$
 $Cons e es \rightarrow \text{case } e \text{ of}$
 $Release_1 \rightarrow f_6 es$
 $| Release_2 \rightarrow f_4 es$
 $| _ \rightarrow f_9 es)$

Figure 7: Result of Distilling Example 1

$$\begin{aligned}
& f_1 \text{ es} \\
& \mathbf{where} \\
& f_1 = \lambda es. \text{Cons (SysState T T) (case es of} \\
& \qquad \qquad \qquad \text{Cons e es} \rightarrow \mathbf{case e of} \\
& \qquad \qquad \qquad \qquad \text{Request}_1 \rightarrow f_2 \text{ es} \\
& \qquad \qquad \qquad \qquad | \text{Request}_2 \rightarrow f_3 \text{ es} \\
& \qquad \qquad \qquad \qquad | _ \qquad \qquad \rightarrow f_1 \text{ es})} \\
& f_2 = \lambda es. \text{Cons (SysState W T) (case es of} \\
& \qquad \qquad \qquad \text{Cons e es} \rightarrow \mathbf{case e of} \\
& \qquad \qquad \qquad \qquad \text{Take}_1 \quad \rightarrow f_4 \text{ es} \\
& \qquad \qquad \qquad \qquad | \text{Request}_2 \rightarrow f_5 \text{ es} \\
& \qquad \qquad \qquad \qquad | _ \qquad \qquad \rightarrow f_2 \text{ es})} \\
& f_3 = \lambda es. \text{Cons (SysState T W) (case es of} \\
& \qquad \qquad \qquad \text{Cons e es} \rightarrow \mathbf{case e of} \\
& \qquad \qquad \qquad \qquad \text{Request}_1 \rightarrow f_5 \text{ es} \\
& \qquad \qquad \qquad \qquad | \text{Take}_2 \quad \rightarrow f_6 \text{ es} \\
& \qquad \qquad \qquad \qquad | _ \qquad \qquad \rightarrow f_3 \text{ es})} \\
& f_4 = \lambda es. \text{Cons (SysState U T) (case es of} \\
& \qquad \qquad \qquad \text{Cons e es} \rightarrow \mathbf{case e of} \\
& \qquad \qquad \qquad \qquad \text{Release}_1 \rightarrow f_1 \text{ es} \\
& \qquad \qquad \qquad \qquad | _ \qquad \qquad \rightarrow f_4 \text{ es})} \\
& f_5 = \lambda es. \text{Cons (SysState W W) (case es of} \\
& \qquad \qquad \qquad \text{Cons e es} \rightarrow \mathbf{case e of} \\
& \qquad \qquad \qquad \qquad _ \rightarrow f_5 \text{ es})} \\
& f_6 = \lambda es. \text{Cons (SysState T U) (case es of} \\
& \qquad \qquad \qquad \text{Cons e es} \rightarrow \mathbf{case e of} \\
& \qquad \qquad \qquad \qquad \text{Release}_2 \rightarrow f_1 \text{ es} \\
& \qquad \qquad \qquad \qquad | _ \qquad \qquad \rightarrow f_6 \text{ es})}
\end{aligned}$$

Figure 8: Result of Distilling Example 2

6 Conclusion and Related Work

In this paper, we have shown how a fold/unfold program transformation technique can be used to verify both safety and liveness properties of reactive systems which have been specified using a functional language. Many corresponding techniques have been developed for verifying temporal properties for logic programs [13, 17, 5, 1, 9]). Some of these techniques have been developed only for safety properties, while others can be used to verify both safety and liveness properties. Due to the use of a different programming paradigm, it is difficult to compare the relative power of these techniques to our own. However, we argue that the use of a more powerful program transformation algorithm will remove more intermediate data structures, and thus be capable of proving more properties directly within the same framework, without the need for making use of external solvers.

Very few techniques have been developed for verifying temporal properties for functional programs other than the work of Lisitsa and Nemytykh [14, 2]. Their approach uses supercompilation [19, 18] as the fold/unfold transformation methodology, where our own approach uses distillation [6, 7]. Since distillation has been shown to be more powerful than positive supercompilation, it follows that we should

$f_1 \text{ es}$
where
 $f_1 = \lambda \text{ es. Cons (SysState T T) (case es of}$
 $\text{Cons e es} \rightarrow \text{case e of}$
 $\text{Request}_1 \rightarrow f_2 \text{ es}$
 $| \text{Request}_2 \rightarrow f_3 \text{ es}$
 $| _ \rightarrow f_1 \text{ es})$
 $f_2 = \lambda \text{ es. Cons (SysState W T) (case es of}$
 $\text{Cons e es} \rightarrow \text{case e of}$
 $\text{Take}_1 \rightarrow f_4 \text{ es}$
 $| \text{Request}_2 \rightarrow f_6 \text{ es}$
 $| _ \rightarrow f_2 \text{ es})$
 $f_3 = \lambda \text{ es. Cons (SysState T W) (case es of}$
 $\text{Cons e es} \rightarrow \text{case e of}$
 $\text{Take}_2 \rightarrow f_5 \text{ es}$
 $| \text{Request}_1 \rightarrow f_7 \text{ es}$
 $| _ \rightarrow f_3 \text{ es})$
 $f_4 = \lambda \text{ es. Cons (SysState U T) (case es of}$
 $\text{Cons e es} \rightarrow \text{case e of}$
 $\text{Release}_1 \rightarrow f_1 \text{ es}$
 $| \text{Request}_2 \rightarrow f_8 \text{ es}$
 $| _ \rightarrow f_4 \text{ es})$
 $f_5 = \lambda \text{ es. Cons (SysState T U) (case es of}$
 $\text{Cons e es} \rightarrow \text{case e of}$
 $\text{Release}_2 \rightarrow f_1 \text{ es}$
 $| \text{Request}_1 \rightarrow f_9 \text{ es}$
 $| _ \rightarrow f_5 \text{ es})$
 $f_6 = \lambda \text{ es. Cons (SysState W W) (case es of}$
 $\text{Cons e es} \rightarrow \text{case e of}$
 $\text{Take}_1 \rightarrow f_8 \text{ es}$
 $| _ \rightarrow f_6 \text{ es})$
 $f_7 = \lambda \text{ es. Cons (SysState W W) (case es of}$
 $\text{Cons e es} \rightarrow \text{case e of}$
 $\text{Take}_2 \rightarrow f_9 \text{ es}$
 $| _ \rightarrow f_7 \text{ es})$
 $f_8 = \lambda \text{ es. Cons (SysState U W) (case es of}$
 $\text{Cons e es} \rightarrow \text{case e of}$
 $\text{Release}_1 \rightarrow f_3 \text{ es}$
 $| _ \rightarrow f_8 \text{ es})$
 $f_9 = \lambda \text{ es. Cons (SysState W U) (case es of}$
 $\text{Cons e es} \rightarrow \text{case e of}$
 $\text{Release}_2 \rightarrow f_2 \text{ es}$
 $| _ \rightarrow f_9 \text{ es})$

Figure 9: Result of Distilling Example 3

be able to verify more properties using our approach. Also, the work of Lisitsa and Nemytykh can verify only safety properties, while our approach can be used to verify both safety and liveness properties.

One other area of work related to our own is the work on using Higher Order Recursion Schemes (HORS) to verify temporal properties of functional programs. HORS are a kind of higher order tree grammar for generating a (potentially infinite) tree and are well-suited to the purpose of verification since they have a decidable mu-calculus model checking problem. Kobayashi [15] first showed how this approach can be used to verify safety properties of higher order functional programs. This approach was then extended to also verify liveness properties by Lester et al. [12]. These approaches have a very bad worst-case time complexity, but techniques have been developed to ameliorate this to a certain extent. It does however appear likely that this approach will be able to verify more properties than our own approach but much less efficiently.

Acknowledgements

This work was supported, in part, by Science Foundation Ireland grant 10/CE/I1855 to Lero - the Irish Software Engineering Research Centre (www.lero.ie), and by the School of Computing, Dublin City University.

References

- [1] Alberto Pettorossi and Maurizio Proietti and Valerio Senni (2009): *Deciding Full Branching Time Logic by Program Transformation*. In: *19th International Symposium on Logic-Based Program Synthesis and Transformation*, pp. 5–21.
- [2] Alexei Lisitsa and Andrei P. Nemytykh (2008): *Reachability Analysis in Verification via Supercompilation*. *International Journal of Foundations of Computer Science* 19(4), pp. 953–969.
- [3] Amir Pnueli and Elad Shahar (1996): *A Platform for Combining Deductive with Algorithmic Verification*. In: *8th International Conference on Computer Aided Verification*, pp. 184–195.
- [4] E.M. Clarke, E.A. Emerson & A.P. Sistla (1986): *Automatic Verification of Finite-State Concurrent Systems Using Temporal Logic Specifications*. *ACM Transactions on Programming Languages and Systems* 8(2), pp. 244–263.
- [5] Fabio Fioravanti and Alberto Pettorossi and Maurizio Proietti (2001): *Verification of Sets of Infinite State Processes Using Program Transformation*. In: *11th International Workshop on Logic Based Program Synthesis and Transformation*, pp. 111–128.
- [6] G.W. Hamilton (2007): *Distillation: Extracting the Essence of Programs*. In: *Proceedings of the ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, pp. 61–70.
- [7] G.W. Hamilton & N.D. Jones (2012): *Distillation With Labelled Transition Systems*. In: *Proceedings of the ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, ACM, pp. 15–24.
- [8] Henny Sipma and Tomás E. Uribe and Zohar Manna (1999): *Deductive Model Checking*. *Formal Methods in System Design* 15(1), pp. 49–74.
- [9] Hirohisa Seki (2011): *Proving Properties of Co-Logic Programs by Unfold/Fold Transformations*. In: *21st International Symposium on Logic-Based Program Synthesis and Transformation*, pp. 205–220.
- [10] L. Lamport (1974): *A New Solution of Dijkstra’s Concurrent Programming Problem*. *Communications of the ACM* 17(8), pp. 453–455.
- [11] Lenore D. Zuck and Amir Pnueli (2004): *Model Checking and Abstraction to the Aid of Parameterized Systems (A Survey)*. *Computer Languages, Systems & Structures* 30(3-4), pp. 139–169.

- [12] Lester, M.M. and Neatherway, R.P. and Ong, C.-H. L. and Ramsay, S.J. (2010): *Model Checking Liveness Properties of Higher-Order Functional Programs*. Unpublished.
- [13] M. Leuschel & T. Massart (1999): *Infinite State Model Checking by Abstract Interpretation and Program Specialisation*. In: *9th International Workshop on Logic Programming Synthesis and Transformation*, pp. 62–81.
- [14] A. Lisitsa & A. Nemytykh (2007): *Verification as a Parameterized Testing (Experiments with the SCP4 Supercompiler)*. *Programming and Computer Software* 33(1), pp. 14–23.
- [15] Naoki Kobayashi (2009): *Types and Higher-Order Recursion Schemes for Verification of Higher-Order Programs*. In: *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL*, pp. 416–428.
- [16] Parosh Aziz Abdulla and Giorgio Delzanno and Noomene Ben Henda and Ahmed Rezine (2009): *Monotonic Abstraction: on Efficient Verification of Parameterized Systems*. *International Journal of Foundations of Computer Science* 20(5), pp. 779–801.
- [17] Abhik Roychoudhury, K. Narayan Kumar, C. R. Ramakrishnan, I. V. Ramakrishnan & Scott A. Smolka (2000): *Verification of Parameterized Systems Using Logic Program Transformations*. In: *Proceedings of the 6th International Conference on Tools and Algorithms for Construction and Analysis of Systems*, pp. 172–187.
- [18] M.H. Sørensen, R. Glück & N.D. Jones (1996): *A Positive Supercompiler*. *Journal of Functional Programming* 6(6), pp. 811–838.
- [19] V.F. Turchin (1986): *The Concept of a Supercompiler*. *ACM Transactions on Programming Languages and Systems* 8(3), pp. 90–121.