Reasoning about Resources in the Embedded Systems Language Hume

Hans-Wolfgang Loidl and Gudmund Grov

School of Mathematical and Computer Sciences
Heriot-Watt University, EH14 4AS Edinburgh, U.K,
{H.W.Loidl,G.Grov}@hw.ac.uk

Abstract. In this paper we present an instrumented program logic for the embedded systems language Hume, suitable to reason about resource consumption. Matching the structure of Hume programs, it integrates two logics, a VDM-style program logic for the functional language and a TLA-style logic for the coordination language of Hume. We present a soundness proof of the program logic, and demonstrate the usability of these logics by proving resource bounds for a Hume program. Both logics, the soundness proof and the example have been fully formalised in the Isabelle/HOL theorem prover.

1 Introduction

Typically resources on embedded systems, such as memory, are very scarce. Thus, an important property of an embedded program is to execute within these limited resources. The Hume embedded systems language achieves a high degree of predictability of resource consumption by defining two language layers: a strict, higher-order functional language at the expression layer, and a restricted language of interacting boxes at the coordination layer. With this design the full computational power of a modern programming language can be used at the expression layer. Where verification and analysis of an expression become intractable, the code can be decomposed into a network of boxes. This language is far more restricted, but easier to analyse and verify.

The different layers require different formalisms for proving (resource) properties. At the expression layer we apply standard techniques from program verification based on a deep embedding of the expression language, which facilitates accurate resource modelling. We develop a Vienna Development Method (VDM)-style [12] program logic to reason about programs. At the coordination layer we find a network of boxes defining their interactions. The properties of interest for this layer can often be described as system invariants. We therefore mechanise a Temporal Logic of Actions (TLA) [14] on this layer. Both layers have been fully formalised in the Isabelle/HOL theorem prover. In this paper, we build a shallow embedding of the coordination layer of Hume on top of an existing, shallow TLA embedding [8] and integrate it with the VDM-style logic.

While these two kinds of logics, and their formalisations, have been studied in isolation, this paper makes three contributions in (i) proving soundness for

the VDM-style program logic in Isabelle/HOL; (ii) combining both logics into an integrated Isabelle/HOL formalisation; and (iii) demonstrating the usability of the combined logics by proving, in Isabelle/HOL, bounded heap consumption of a simple Hume program.

2   The Hume Language

Hume [9] is designed as a layered language where the *coordination layer* is used to construct reactive systems using a finite-state-automata based notation, representing a static system of interconnecting *boxes*; while the *expression layer* describes the behaviour of a box using a strict, higher-order, purely functional rule-based notation. A central design goal of Hume is predictability of resource consumption, so that each expression-layer program can execute within bounded time and space constraints. Thus, we are mainly interested in proving resource bounds for Hume programs.

The *expression layer* of Hume corresponds to a strict, higher-order language with general algebraic datatypes. To accurately model resource consumption, we formalise the expression layer of Hume in the form that is used as an intermediate language in the compiler, distinguishing between first- and higher-order function calls for performance reasons. This language makes some structural restrictions on the source code, most notably it is in let-normal-form, and uses general algebraic data-types:

$$\text{Patt} \ni p ::= x \mid v \mid c\ x_1 \ldots x_n \mid \_$$
$$\text{Expr} \ni e ::= x\ \mid v \mid c\ x_1 \ldots x_n \mid f\ x_1 \ldots x_n \mid x\ x_1 \ldots x_n \mid \texttt{if}\ x\ \texttt{then}\ e_1\ \texttt{else}\ e_2 \mid$$
$$x_1 \oplus x_2 \mid \texttt{let}\ x\ = e_1\ \texttt{in}\ e_2 \mid \texttt{case}\ x\ \texttt{of}\ p_1 \to e_1\ \texttt{otherwise}\ e_2$$

A pattern p is either a variable $x \in \text{Var}$, a value $v \in \text{Val}$, a constructor application $c\ x_1 \ldots x_n$ ($c \in \text{Constr}$) or a wildcard $\_$. An expression e is either a variable x, a value v, a constructor application $c\ x_1 \ldots x_n$, a first-order function call of f or a higher-order function call of x with arguments $x_1 \ldots x_n$, a conditional, a binary primitive operation $\oplus$, a let-expression, or a case expression. The latter is a one-step matching expression, with a default branch to be used if the match was unsuccessful. The entire program is represented by a table funTab, which maps a function name (f) to its formal arguments ($args_f$) and to its body ($body_f$).

The *coordination layer* of Hume comprises of a set of concurrent, asynchronous boxes, scheduled in an alternating sequence of execute and super-step phases. In the *execute* phase any box with sufficient available inputs will be executed. In the following *super-step* phase, the results of the boxes will be copied to single-buffer wires, provided they are free. Boxes are defined through pattern-matching rules, of the form $p \to e$, using the syntax of the expression language above. The semantics of these rules is similar to the one of expression-layer pattern matching, but additionally accounts for possibly missing input values.

To illustrate, Figure 1, shows a configuration of two connected boxes, `copy` followed by `fuse`, which will be our running example in this paper: Here, both
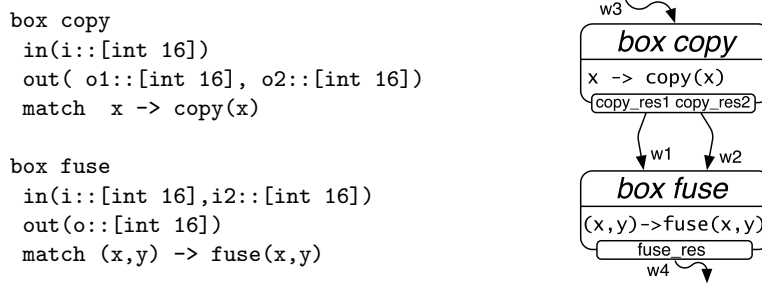
```
box copy
 in(i::[int 16])
 out( o1::[int 16], o2::[int 16])
 match  x -> copy(x)

box fuse
 in(i::[int 16],i2::[int 16])
 out(o::[int 16])
 match (x,y) -> fuse(x,y)
```


Fig. 1. The Running Hume Example

$$
\begin{array}{llll}
l,a,r & \in & \mathrm{Locn} & = \mathbb{N} \uplus \{\mathrm{nil}\} \\
h,h' & \in & \mathrm{Heap} & = \mathrm{Locn} \rightsquigarrow_f \mathrm{Val} \\
v & \in & \mathrm{Val} & = \{\bot\} \uplus \mathbb{Z} \uplus \mathbb{B} \uplus \mathrm{Locn} \uplus (\mathrm{Constr}, \mathrm{Locn}^*) \uplus (\mathrm{Funs}, \mathbb{N}, \mathrm{Locn}^*) \\
P,Q,A & \in & \mathrm{Assn} & = \mathrm{Env} \Rightarrow \mathrm{Heap} \Rightarrow \mathrm{Heap} \Rightarrow \mathrm{Val} \Rightarrow \mathrm{Resources} \Rightarrow \mathbb{B} \\
p & \in & \mathrm{Resources} & = (clock :: \mathbb{Z},\ callc :: \mathbb{Z}, maxstack :: \mathbb{Z})
\end{array}
$$

$$
\begin{array}{lll}
E,E' & \in \mathrm{Env} & = \mathrm{Var} \Rightarrow \mathrm{Val} \\
G & \in \mathrm{Ctxt} & = \mathcal{P}\{(\mathrm{Expr}, \mathrm{Assn})\}
\end{array}
$$

Fig. 2. Basic Domains

boxes work on lists of 16-bit integers as input value, and calls two respective functions which we will return to later. In general, the body of a box is comprised of a sequence of rules, with the left hand side being matched against the input data, and first successful match enabling execution of its right hand side. The two output wires of `copy` are the inputs of `fuse`, and we will in this discussion ignore the source input of `copy` and the destination output of `fuse`.

3   A VDM-style Program Logic for the Expression Layer

3.1   Operational Semantics

Figure 2 summarises the domains used in the formalisation. The sets of identifiers for variables ($x \in \mathrm{Var}$), functions ($f \in \mathrm{Funs}$), and data constructors ($c \in \mathrm{Constr}$) are disjoint. We represent heaps ($h \in \mathrm{Heap}$) by finite mappings from locations to values (written $\rightsquigarrow_f$), and environments ($E \in \mathrm{Env}$) by total functions of variables to values (written $\Rightarrow$). A location $l \in \mathrm{Locn}$ is either a natural number (representing an address in the heap), or the constant nil. We write Locn^* for the domain of sequences of elements in Locn, $E \star \vec{xs}$ for mapping E over the sequence $\vec{xs}$, prefix # for the length of a sequence and ++ for append of sequences. We model heap usage of h as the size of its domain, i.e. $|\operatorname{dom} h|$. Other resources, such as time, are modelled in the *resource vector* $p \in \mathrm{Resources}$ and $\smile$ combines these. We obtain costs for basic operations from a parameterised table $\mathcal{R}$. The operation freshloc has the property: freshloc $s \notin s$, for all s.

$$\frac{}{E, h \vdash v \Downarrow_1 (v, h, \mathcal{R}^{\mathtt{const}})} \ (\textsc{Value}) \quad \frac{E\ x = v}{E, h \vdash x \Downarrow_1 (v, h, \mathcal{R}^{\mathtt{var}})} \ (\textsc{Var})$$

$$\frac{\begin{array}{c} l = \text{freshloc (dom } h) \\ \vec{rs} = E \star \vec{xs} \qquad h' = h(l \mapsto (c,\ \vec{rs})) \end{array}}{E, h \vdash c\ \vec{xs} \Downarrow_1 (l, h', \mathcal{R}^{\mathtt{constr}\ \#\vec{xs}})} \ (\textsc{Constr}) \quad \frac{\begin{array}{c} v_1 = E\ x_1 \qquad v_2 = E\ x_2 \\ v = v_1 \oplus v_2 \end{array}}{E, h \vdash x_1 \oplus x_2 \Downarrow_1 (v, h, \mathcal{R}^{\oplus})} \ (\textsc{PrimBin})$$

$$\frac{E\ x = \text{true} \qquad E, h \vdash e_1 \Downarrow_n (v, h', p')}{E, h \vdash \mathtt{if}\ x\ \mathtt{then}\ e_1\ \mathtt{else}\ e_2 \Downarrow_{n+1} (v, h', p' \smile \mathcal{R}^{\mathtt{if\ true}})} \ (\textsc{IfTrue})$$

$$\frac{E\ x = \text{false} \qquad E, h \vdash e_2 \Downarrow_n (v, h', p')}{E, h \vdash \mathtt{if}\ x\ \mathtt{then}\ e_1\ \mathtt{else}\ e_2 \Downarrow_{n+1} (v, h', p' \smile \mathcal{R}^{\mathtt{if\ false}})} \ (\textsc{IfFalse})$$

$$\frac{E, h \vdash e_1 \Downarrow_m (v, h', p') \qquad E(x := v), h' \vdash e_2 \Downarrow_n (v'', h'', p'')}{E, h \vdash \mathtt{let}\ x\ =\ e_1\ \mathtt{in}\ e_2 \Downarrow_{m+n} (v'', h'', p' \smile p'' \smile \mathcal{R}^{\mathtt{let}})} \ (\textsc{Let})$$

$$\frac{l = E\ x \quad \mathtt{MATCH}\ E, h \vdash p\ \mathtt{at}\ l \Downarrow (E', v') \quad v' \neq \bot \quad E', h \vdash e_1 \Downarrow_n (v, h', p')}{E, h \vdash \mathtt{case}\ x\ \mathtt{of}\ p \to e_1\ \mathtt{otherwise}\ e_2 \Downarrow_{n+1} (v, h', p' \smile \mathcal{R}^{\mathtt{case\ true}})} \\ (\textsc{CaseTrue})$$

$$\frac{l = E\ x \quad \mathtt{MATCH}\ E, h \vdash p\ \mathtt{at}\ l \Downarrow (E', v') \quad v' = \bot \quad E, h \vdash e_2 \Downarrow_n (v, h', p')}{E, h \vdash \mathtt{case}\ x\ \mathtt{of}\ p \to e_1\ \mathtt{otherwise}\ e_2 \Downarrow_{n+1} (v, h', p' \smile \mathcal{R}^{\mathtt{case\ false}})} \\ (\textsc{CaseFalse})$$

$$\frac{E' = E(args_f := E \star \vec{xs}) \qquad E', h \vdash body_f \Downarrow_n (v, h', p)}{E, h \vdash f\ \vec{xs} \Downarrow_{n+1} (v, h', p \smile \mathcal{R}^{f\ \#\vec{xs}})} \ (\textsc{CallFun})$$

$$\frac{\begin{array}{c} l = E\ x \qquad \text{Some } (f,\ i,\ \vec{rs}) = h\ l \qquad i + \#\vec{xs} < \#args_f \\ l' = \text{freshloc (dom } h) \qquad h' = h(l' \mapsto (f,\ (i + \#\vec{xs}),\ (\vec{rs} \mathbin{+\!\!+} (E \star \vec{xs})))) \end{array}}{E, h \vdash x\ \vec{xs} \Downarrow_1 (l', h', \mathcal{R}^{\mathtt{ap\ false}\ \#\vec{xs}})} \\ (\textsc{CallVarUnderApp})$$

$$\frac{\begin{array}{c} l = E\ x \qquad \text{Some } (f,\ i,\ \vec{rs}) = h\ l \qquad i + \#\vec{xs} = \#args_f \\ E' = E(args_f := \vec{rs} \mathbin{+\!\!+} (E \star \vec{xs})) \qquad E', h \vdash body_f \Downarrow_n (v, h', p) \end{array}}{E, h \vdash x\ \vec{xs} \Downarrow_{n+1} (v, h', p \smile \mathcal{R}^{\mathtt{ap\ true}\ \#\vec{xs}})} \\ (\textsc{CallVarExact})$$

Fig. 3. Operational Semantics of Expression-layer Hume

A judgement of the big-step operational semantics, shown in Figure 3, has the form $E, h \vdash e \Downarrow_m (v, h', p)$ and is read as follows: given a variable environment, E, and a heap, h, e evaluates in m steps to the value v, yielding the modified heap, h', and consumes resources p. A value v is itself the result (VALUE). For a variable x a lookup is performed in E (VAR). For a constructor application $c \, \vec{xs}$ the heap is extended with a fresh location, mapped to c and its argument values (CONSTR). For primitive operations a shallow embedding is used so that the operator $\oplus$ can be directly applied to the values of its variables (PRIMBIN). A conditional returns the value of the then branch e_1 or the else branch e_2, depending on the value of x. A let binds the result of e_1 to the variable x and returns the value of the body e_2 (LET). The judgement of the pattern-matching semantics, MATCH $E, h \vdash p$ at $l \Downarrow$ (E', l'), matches the value at location l against pattern p, returning a modified environment and l on success, or the same environment and $\perp$ on failure. A case is a one-step matching expression, that matches a variable x against a pattern p; if successful its value is that of e_1 (CASETRUE), if unsuccessful its value is that of e_2 (CASEFALSE). In first-order function calls, we bind the values in the arguments to the formal parameters ($args_f$), and use the modified environment to evaluate the body ($body_f$) of the function f (CALLFUN). For higher-order function calls, with function closures represented as triples $(f, i, \vec{rs})$ of function name, number of arguments and their values, we have to distinguish two cases. If we have fewer than the required number of arguments, the result will be a new closure consisting of f, the number of arguments supplied so far and their values. If we have precisely the number of arguments that the function requires, the values in the closure and the those of the arguments ($\vec{xs}$) are bound to the formal parameters ($args_f$) of the function, before the function body is evaluated using these new bindings (CALLVAREXACT).

3.2   Program Logic

We use a VDM-style of assertions [12] to model properties of programs. This means that rather than using the more common Hoare-triples of pre-condition, program and post-condition, we use a judgement of the following form, $G \rhd e :$ P, meaning that expression e fulfils the assertion P in a context G. As shown in Figure 2, an *assertion* is a predicate over the components of the operational semantics, and a context is a set of pairs of program expression $e \in$ Expr and assertion $P \in$ Assn. See [13] for a comparison of VDM- and Hoare-style logics.

The rules in Figure 4 define the program logic for expression-layer Hume. A value w is the result value, and the heap is unchanged (VDMVALUE). A variable x requires a lookup in the environment, and the heap is unchanged (VDMVAR). In a constructor application, the newly allocated location is existentially quantified, and the heap is updated with a binding to this location (VDMCONSTR). The two possible control flows in the conditional are encoded in the logic as implications, based on the boolean contents of variable x (VDMIF). In the let rule (VDMLET), the intermediate value, heap and resources are existentially quantified, and the environment for executing e_2 is updated accordingly. The control flows in the case construct are encoded as implications, based on the result from the pattern

$$\frac{}{G \ \triangleright \ w \ : \ \lambda E \ h \ h' \ v \ p. \ v = w \ \wedge \ h = h' \ \wedge \ p = \mathcal{R}^{\text{const}}} \quad (\textsc{vdmValue})$$

$$\frac{}{G \ \triangleright \ x \ : \ \lambda E \ h \ h' \ v \ p. \ v = E \ x \ \wedge \ h = h' \ \wedge \ p = \mathcal{R}^{\text{var}}} \quad (\textsc{vdmVar})$$

$$\frac{}{\begin{array}{c} G \ \triangleright \ c \ \vec{xs} \ : \ \lambda E \ h \ h' \ v \ p. \ \exists l \ \vec{rs}. \ l = \text{freshloc}(\text{dom} \ h) \ \wedge \ \vec{rs} = E \star \vec{xs} \ \wedge \\ v = l \ \wedge \ h' = h(l \mapsto (c, \ \vec{rs})) \ \wedge p = \mathcal{R}^{\text{constr} \ \#\vec{xs}} \end{array}} \quad (\textsc{vdmConstr})$$

$$\frac{}{\begin{array}{c} G \ \triangleright \ x_1 \oplus x_2 \ : \lambda E \ h \ h' \ v \ p. \ \exists v_1 \ v_2. \ E \ x_1 = v_1 \ \wedge \ E \ x_2 = v_2 \ \wedge \\ v = v_1 \oplus v_2 \ \wedge \ h = h' \ \wedge \ p = \mathcal{R}^{\oplus} \end{array}} \quad (\textsc{vdmPrimBin})$$

$$\frac{G \ \triangleright \ e_1 \ : P_1 \qquad G \ \triangleright \ e_2 \ : P_2}{\begin{array}{c} G \ \triangleright \ \text{if } x \text{ then } e_1 \text{ else } e_2 \ : \ \lambda E \ h \ h' \ v \ p. \\ (E \ x = \text{true} \ \longrightarrow \ \exists p'. \ P_1 \ E \ h \ h' \ v \ p' \ \wedge \ p = p' \smile \mathcal{R}^{\text{if true}}) \wedge \\ (E \ x = \text{false} \ \longrightarrow \ \exists p'. \ P_2 \ E \ h \ h' \ v \ p' \ \wedge \ p = p' \smile \mathcal{R}^{\text{if false}}) \end{array}} \quad (\textsc{vdmIf})$$

$$\frac{G \ \triangleright \ e_1 \ : P_1 \qquad G \ \triangleright \ e_2 \ : P_2}{\begin{array}{c} G \ \triangleright \text{let } x \ = \ e_1 \text{ in } e_2 \ : \ \lambda E \ h \ h' \ v \ p. \ \exists v' \ h'' \ p' \ p''. \\ (P_1 \ E \ h \ h'' \ v' \ p' \ \wedge \ P_2 \ E(x := v') \ h'' \ h' \ v \ p'' \ \wedge \ p = p' \smile p'' \smile \mathcal{R}^{\text{let}}) \end{array}} \quad (\textsc{vdmLet})$$

$$\frac{G \ \triangleright \ e_1 \ : \ P_1 \qquad G \ \triangleright \ e_2 \ : \ P_2}{\begin{array}{c} G \ \triangleright \ \text{case } x \text{ of } p_1 \rightarrow e_1 \text{ otherwise } e_2 \ : \ \lambda E \ h \ h' \ v \ p. \ \forall l. \ l = E \ x \ \longrightarrow \ \forall E' \ v'. \\ \text{MATCH } E, \ h \vdash \ p_1 \text{ at } l \ \Downarrow \ (E', v') \ \longrightarrow \\ ((v' = \bot \ \longrightarrow \ \exists p'. \ P_2 \ E \ h \ h' \ v \ p' \ \wedge \ p = p' \smile \mathcal{R}^{\text{case false}}) \ \wedge \\ (v' \neq \bot \ \longrightarrow \ \exists p'. \ P_1 \ E' \ h \ h' \ v \ p' \ \wedge \ p = p' \smile \mathcal{R}^{\text{case true}})) \end{array}} \quad (\textsc{vdmCase})$$

$$\frac{\begin{array}{c} \{(f \ \vec{xs}, P)\} \cup G \ \triangleright \ body_f \ : \ \lambda E \ h \ h' \ v \ p. \ \forall E'. \ E = E'(args_f := E' \star \vec{xs}) \ \longrightarrow \\ P \ E' \ h \ h' \ v \ (p \smile \mathcal{R}^{f \ \#\vec{xs}}) \end{array}}{G \ \triangleright \ f \ \vec{xs} \ : \ P} \quad (\textsc{vdmCallFun})$$

$$\frac{\begin{array}{c} \forall E \ h \ h' \ v. \ \Phi_{x,\vec{xs}} \ E \ h \ h' \ v \ \longrightarrow \ P \ E \ h \ h' \ v \ \mathcal{R}^{\text{ap false} \ \#\vec{xs}} \\ (\forall f. \ \{(x \ \vec{xs}, P)\} \cup G \ \triangleright \ body_f \ : \ \lambda E \ h \ h' \ v \ p. \ \forall E'. \ (\exists \vec{rs}. \ \Psi_{x,\vec{xs}} \ E' \ h \ h' \ v \ f \ \vec{rs} \ \wedge \\ E = E'(args_f := \vec{rs} ++ (E' \star \vec{xs})) \ \longrightarrow \ P \ E' \ h \ h' \ v \ (p \smile \mathcal{R}^{\text{ap true} \ \#\vec{xs}})) \end{array}}{G \ \triangleright \ x \ \vec{xs} \ : \ P} \quad (\textsc{vdmCallVar})$$

$$\frac{(e, P) \in G}{G \ \triangleright \ e \ : \ P} \ (\textsc{vdmAx}) \qquad \frac{\forall E \ h \ h' \ v \ p. \ P \ E \ h \ h' \ v \ p \ \longrightarrow \ Q \ E \ h \ h' \ v \ p \qquad G \ \triangleright \ e \ : \ P}{G \ \triangleright \ e \ : \ Q} \ (\textsc{vdmConseq})$$

$$\begin{array}{c} \Phi_{x,\vec{xs}} \equiv \lambda E \ h \ h' \ v. \ \exists l \ l' \ f \ i \ \vec{rs}. \ E \ x = l \ \wedge \ h \ l = \text{Some}(f, \ i, \ \vec{rs}) \ \wedge \ i + \#\vec{xs} < \#args_f \ \wedge \\ l' = \text{freshloc} \ (\text{dom} \ h) \ \wedge \ v = l' \ \wedge \ h' = h(l' \mapsto (f, \ (i + \#\vec{xs}), \ (\vec{rs} ++ (E \star \vec{xs})))) \\ \Psi_{x,\vec{xs}} \equiv \lambda E' \ h \ h' \ v \ f \ \vec{rs}. \ \exists l \ i \ . \ E' \ x = l \ \wedge \ h \ l = \text{Some} \ (f, \ i, \ \vec{rs}) \ \wedge \ i + \#\vec{xs} = \#args_f \end{array}$$

Fig. 4. Program Logic for Expression-layer Hume

match (VDMCASE). In a first order call (VDMCALLFUN), the assertion to be proven for f is added with the call into the context, and the function body has to fulfil the assertion, after modifying environment and resources. The higher-order function call (VDMCALLVAR) directly encodes the preconditions of the two cases of under application (Φ) and of exact application (Ψ) from the operational semantics. Thus, the definition of Φ reads as follows: for the pre-state E, h we find in variable x a closure with i arguments, and the total number of arguments is smaller than the function's arity $(i + \#\vec{xs} < \#args_f)$. In this case, the result state v, h', p is constructed such that the result value is a new closure, containing all arguments, and the result heap is updated accordingly. The definition of Ψ reads as follows: For the pre-state E', h we find in variable x a closure with i arguments, and the total number of arguments matches the function's arity $(i + \#\vec{xs} = \#args_f)$. In this case, the second premise in rule VDMCALLVAR demands that the body of f $(body_f)$ fulfils the assertion P, with the environment adjusted for parameter passing and a modified resource vector to account for the costs of the function call. Note that Ψ has to be parameterised over the function f and the argument values $\vec{rs}$ captured in its closure, so that the VDMCALLVAR rule can quantify over f and $\vec{rs}$. In particular, the quantification over f scopes over the entire second judgement in the premise of VDMCALLVAR, because it is needed to retrieve the function's body via $body_f$. Finally, the quantification over $\vec{rs}$ scopes over the entire pre-condition inside the assertion. The environment E is constructed out of E' by binding the formal parameters to the values in the closure $(\vec{rs})$ and those retrieved from the arguments $\vec{xs}$. The final two rules, VDMAX and VDMCONSEQ, are the standard rules for using an axiom, present in the context, and for logical consequence in the meta-language.

3.3   Soundness

We define (relativised) semantic validity of an assertion for an expression in a context as follows.

Definition 1 (validity). *Assertion P is valid for expression e ($\models_n e : P$) iff*

$$\forall m \leq n. \ \forall E \ h \ h' \ v \ p. \ (E, \ h \ \vdash \ e \ \Downarrow_m \ (v, h', p)) \ \longrightarrow \ P \ E \ h \ h' \ v \ p$$

Assertion P is valid for expression e in context G ($G \models e : P$) iff

$$\forall n. \ (\forall (e', P') \ \in \ G. \ \models_n \ e' \ : \ P') \ \longrightarrow \ \models_n \ e \ : \ P$$

Based on these simple definitions of validity, exploiting the shallow nature of our embedding of assertions, the soundness theorem can be stated as follows.

Theorem 1 (soundness). *For all contexts G, expressions e, assertions P*

$$G \triangleright e : P \Longrightarrow G \models e : P$$

Proof structure: By induction over the rules of the program logic. In the case of function calls, induction over the index n in the semantics relation[1]. $\square$

Our approach to proving *mutually recursive functions*, building on [1], is to define a predicate *goodContext*. In the first-order case, it requires context elements to be function calls associated to entries in the function specification table $\mathcal{F}$, and, informally, the context must be powerful enough to prove all of its assertions. This means that it has to encode information about all functions called in the body of the function under consideration. Note the universal quantification over the arguments $\vec{ys}$ of the function. This allows us to adapt the function arguments to the concrete values provided at the call site.

Definition 2. *A context G is called a good context w.r.t. specification table $\mathcal{F}$, written goodContext $\mathcal{F}$ G, iff*

$$
\begin{aligned}
&\forall e\ P\ .\ (e, P) \in G\ \longrightarrow \\
&\quad \exists f\ \vec{xs}.\ e = f\ \vec{xs}\ \wedge\ P = \mathcal{F}\ f\ \vec{xs}\ \wedge \\
&\qquad \forall \vec{ys}.\ G \rhd body_f\ : \\
&\qquad\quad \lambda E\ h\ h'\ v\ p.\ \forall E'.\ E = E'(args_f := E' \star \vec{ys})\ \longrightarrow \\
&\qquad\qquad (\mathcal{F}\ f\ \vec{ys})\ E'\ h\ h'\ v\ (p \smile \mathcal{R}^{f\ \#\vec{ys}})
\end{aligned}
$$

$$
\frac{G\ \rhd\ e : P}{G \cup G'\ \rhd\ e : P}\ (\textsc{ctxtWeak}) \qquad \frac{G\ \rhd\ e' : P' \qquad (\{(e', P')\} \cup G)\ \rhd\ e : P}{G \rhd\ e : P}\ (\textsc{cut})
$$

$$
\frac{\text{finite } G \qquad \mid G \mid\, = n \qquad goodContext\ \mathcal{F}\ G \qquad (e, P) \in G}{\emptyset\ \rhd\ e : P}\ (\textsc{mutrec})
$$

$$
\frac{goodContext\ \mathcal{F}\ G \qquad \text{finite } G \qquad (f\ \vec{xs}, \mathcal{F}\ f\ \vec{xs})\ \in G}{\emptyset\ \rhd\ f\ \vec{ys}\ :\ \mathcal{F}\ f\ \vec{ys}}\ (\textsc{vdmAdapt})
$$

Fig. 5. Admissible Rules

Figure 5 shows a set of *admissible rules* that are useful in proving concrete program properties. The rules ctxtWeak and cut are proven by induction over the derivation of $G \rhd\ e : P$. Rule mutrec is proven by induction over the size of G. Finally, vdmAdapt, which is used when proving a property on a function call, follows directly from mutrec. Note that vdmAdapt reduces the proof to one of the above good context predicate that has to be constructed for the function under consideration. Thus, proving *goodContext $\mathcal{F}$ G*, for a specially constructed context, becomes the main step in proving a property for (mutually) recursive functions. Section 5 gives an example.

---

[1] This is theorem `soundness` in the Isabelle/HOL theory `GRUMPY.thy`.

4   Incorporating the Coordination Layer using TLA

While VDM is suitable for the expression layer, it is not a natural way of reasoning about properties of the coordination layer. This is a stateful layer, where boxes communicate over single-buffered wires. Here, we are interested in invariance properties for the overall program, expressed using the wires. For such networks and properties, a logic such as TLA [14], originally developed to reason about concurrent systems, is more suitable, and has been formalised in Isabelle/HOL [8]. A TLA specification contains a set of variables, and invariants are predicates over these variables, preceded by the temporal $\square$ (always) operator. The behaviour of the system is written in what is known as the *action level*, which logically is a predicate over two states: a before state and an after state, separated by priming all variables in the after state. Let $\mathcal{N}$ be such an action predicate, for a specification with variables v. Moreover, let I be a state predicate (i.e. a predicate over a single state), which represents the initial state of all variables. A specification of this system is then written $I \wedge \square[\mathcal{N}]_v$, meaning that I holds initially and every step henceforth is either a step by $\mathcal{N}$ or a stuttering step, leaving v unchanged (such stuttering a step enables internal steps which are useful when refining a program which is not considered further here). As both program and properties are expressed in the same logic, the invariant P is proven by proving the implication $I \wedge \square[\mathcal{N}]_v \longrightarrow \square P$, i.e. P holds initially, and holds for all steps $\mathcal{N}$ can make. The key to proving this is the following Inv rule of TLA [14]:

$$\frac{I \longrightarrow P \qquad P \wedge \mathcal{N} \longrightarrow P' \qquad P \wedge v' = v \longrightarrow P'}{I \wedge \square[\mathcal{N}]_v \longrightarrow \square P} \qquad \text{(Inv)}$$

TLA is mechanised in Isabelle/HOL using possible world semantics [6]. This embedding is shallow, which enables direct use of existing Isabelle/HOL tools and theorems. Moreover, the Hume embedding is also shallow, which simplifies the expression layer integration.

Hume boxes are executed in a two-step lock step algorithm [7]: in the first execute phase all runnable boxes are executed sequentially; this is followed by a super-step, which copies box outputs to wires. This is embedded by two (enumeration) meta-variables: a scheduler s and a program counter pc.

The scheduling is formalised by an action S, which updates the scheduling variable s. The scheduler depends on whether pc' equals the (defined) "last box" in the enumeration type.

Definition 3. *By giving the value of this "last box" predicate as argument, the scheduler is defined as follows:*

$$S \; last\_box \equiv s' = (\texttt{if } s = Execute \wedge last\_box \texttt{ then } Super \texttt{ else } Execute)$$

Each Hume wire is represented by a variable, and to achieve the lock-step scheduling, each box has one result buffer variable for each output. This is illustrated in Figure 1, which shows the example discussed in Section 5. In addition,

each box has a state/mode variable (e.g. *copy_st*) of the enumeration type consisting of *Runnable*, *Blocked*, and *Matchfail*: a *Runnable* box can be executed; *Matchfail* denotes that the box has failed in matching the inputs and cannot be executed; a *Blocked* box has failed asserting the output buffer with the output wires, and cannot be executed. All the result buffers and wires are of type Val. In addition each program has one heap $h \in$ Heap and a $p \in$ Resources which specifies resource properties. Every box has one action for each phase.

In the super-step, box outputs are checked, and if successful the result buffer is copied to the wires. For heap structures, this is a shallow copy of the references on the stack, thus the heap is left unchanged. In the execute phase, only the current box (identified by *pc*) is executed in a given step. Moreover, it is only executed when it is not *Blocked*. In a box execution, there is first a check that there is a match that will succeed (a box may not be total). If this check succeeds, $E, h \vdash e \Downarrow (v, h', p)$ is used to represent the execution of the expression layer e, under the environment E, with a correct variable (Var) to value (Val) mapping and the current heap h.

To enable usage of the inductively defined judgement $E, h \vdash e \Downarrow (v, h', p)$ within the TLA-embedded expression layer, it must be seen as a function, accepting the old heap h, environment E and expression e as input, and returning the new heap h', the return value v and resources used p.

Definition 4. *We define the* exe *function as follows:*

$$((v, h', p) = exe\ E\ h\ e) \equiv (E,\ h\ \vdash\ e\ \Downarrow\ (v, h', p))$$

The *exe* function is then used to update the result buffer, heap h and resource variable p in the TLA representation. The proof of TLA invariants of the coordination layer relies on pre- and post-conditions of the *exe* function. To achieve this, the *exe* function must be linked with the VDM logic for proofs of properties of the form $E, h \vdash e \Downarrow (v, h', p)$. This is achieved by the following theorem, which is the key for the integration between the TLA and VDM logic:

Theorem 2 (vdmexe[2]). $[\![((v, h', p) = exe\ E\ h\ e); \emptyset \rhd e : A]\!] \Longrightarrow A\ E\ h\ h'\ v\ p$

Proof. The assumption applied to the *soundness* theorem gives $\emptyset \vDash e : A$. By the definition of *validity in context*, this gives $\vDash \emptyset \longrightarrow (\vDash e : A)$ which gives (G1) $\vDash e : A$. Moreover, the definition of *exe* and the assumption of *vdmexe*, implies $E, h \vdash e \Downarrow (v, h', p)$. The goal then follows from unfolding (G1), by using the definition of *validity*, followed by an application of it. □

This integration works directly due to the shallow TLA embedding and the shallow assertion language for a VDM property. VDM properties are then explored within TLA, since the *vdmexe* theorem turns the VDM property into a HOL predicate, enabling the VDM proof to be independent of TLA. We will now illustrate how the VDM and TLA logics are combined in the proof of a Hume invariant $\Box P$. This approach will be known as the "*standard Hume/TLA invariant proof structure.*"

---

[2] This is theorem `vdmexe` in `Integrated.thy`.

Definition 5 (the standard Hume/TLA invariant proof structure). *As a first step the* Inv *rule above is applied. The base (initial state) $I \Rightarrow P$ and "unchanged" $P \land v' = v \Rightarrow P'$ cases are trivial and not discussed further.*

In the main $P \land \mathcal{N} \Rightarrow P'$ case, $\mathcal{N}$ is a conjunction of the scheduler S and all the box actions. It starts by case-analysis on s, creating two cases. The first case, $s = Super$ is achieved purely within the TLA logic. The second case, $s = Execute$ is followed by a case analysis on pc. *If P' depends on the result of execution* pc, *then it is followed by a case-split on $pc\_st = Blocked$ (the box state/mode variable). Let e be the expression layer of this box. If the check for a match succeeds, then the proof of P' depends on a sufficiently strong A such that $\emptyset \triangleright e : A$ can be proved by the VDM-logic. The vdmexe theorem is then applied to enable the use of A in the main TLA proof of P'.*

5   Example — a Verified List-Copy Configuration

To illustrate the integrated reasoning infrastructure, we will verify a resource property over the heap, for our running example shown in Figure 1[3]. The *copy* box performs a list-copy over its input list, producing two identical lists; the *fuse* box takes these two lists and combines them into one list, using only the smaller of the two elements from the input lists. We first define these two functions in let-normal-form:

```
copy [x] ≡ case x of
   (CONS [h,t]) → let y = copy [t] in CONS [h,y]   otherwise   NIL []
fuse [x₀,x₁] ≡ case x₀ of (CONS [h₀,t₀]) →   case x₁ of   (CONS [h₁,t₁]) →
   let y = fuse [t₀,t₁] in let b = (h₀ < h₁) in
      if b then CONS [h₀,y]  else CONS [h₁,y]
   otherwise   NIL []   otherwise   NIL []
```

Here, the *fuse* box accepts two inputs (x, y), binding them in the environment, and the *copy* box returns two values. This is captured by the following interface: $(fuse\_res', h', r') = exe\ \emptyset(x := w1, y := w2)\ h\ (fuse\ [x, y])$. For *copy* we model the fact that the input from $w3$ is forwarded directly to its first output as follows: $(copy\_res1', (copy\_res2', h', r')) = (w3, exe\ \emptyset(x := w3)\ h\ (copy\ [x]))$. The super-step is similar to the one already described, although the copy box checks and updates two wires.

In this example the box network is open, with input wire $w3$ and output wire $w4$. Therefore we have to extend $\mathcal{N}$ with an additional conjunct *env*, which models the environment. Both $w3$ and $w4$ are assumed to be updated by other boxes (or streams). In the coordination-layer proof below, we model a single, non-overlapping evaluation where wire $w3$ initially contains a list structure (*mList*) and is never updated, while wire $w4$ consumes the result.

---

[3] This example is theory `ListCopy.thy` of the Isabelle/HOL embedding.

We specify the layout of a list structure by the following inductive definition *mList*. Informally, $(n, a, U, h) \in mList$ means that at address a in heap h we find a list structure with n CONS and 1 NIL cells, covering the addresses in U.

$$\frac{h\ a = \text{Some (NIL, [])}}{(0, a, \{a\}, h) \in mList}\ \text{mListNIL} \qquad \frac{h\ a = \text{Some (CONS, } [r, r'])\quad a \notin U \quad (n, r', U, h) \in mList}{(n + 1, a, U \cup \{a\}, h) \in mList}\ \text{mListCONS}$$

The following definition is useful to express that address a contains a list of length n: *list h a n* $\equiv \exists\ U.\ (n, a, U, h) \in mList$

We are interested in the heap consumption, specified as a relation between the pre-heap h and the post-heap h'. We define these resource properties as entries in the *function specification table* $\mathcal{F}$. The specification for *copy* states, that provided the argument x points to a list structure of length n in the pre-heap h, the size of the post-heap h' is at most the size of the pre-heap plus $n+1$. Thus, the heap consumption of the *copy* function is $n + 1$. Furthermore, the result of the function (in v) is a list of length n at location r' in the post-heap h'. The remaining clauses assure that the input data structure is not modified. The clause $h \sim_{(\text{dom}\ h') - U'} h'$ states that the heaps h and h' contain the same values at addresses $(\text{dom}\ h') - U'$, i.e. only the values in U' are modified. The clause $U' \cap (\text{dom}\ h) = \emptyset$ states that the copy of the list does not overlap with the input list, and the $(n, r, U, h') \in mList$ clause states that the structure of the input list is indeed unchanged. Analogously, the specification for *fuse* states, that provided the arguments x_0 and x_1 point to list structures, of the same length n_0 and n_1, at locations r_0 and r_1 in the pre-heap h, the size of the post-heap h' is at most the size of the pre-heap plus $n_0 + 1$. The result of the function is a list of length n_0 at location r' in the post-heap.

$\mathcal{F}\ copy\ [x] \equiv \lambda E\ h\ h'\ v\ p.\ \forall n\ r\ U.$
   $E\ x\ =\ r\ \wedge\ (n, r, U, h) \in mList \longrightarrow$
      $\mid \text{dom}\ h' \mid \leq \mid \text{dom}\ h \mid +n + 1\ \wedge (\exists r'\ U'.\ v = r'\ \wedge\ (n, r', U', h') \in mList\ \wedge$
      $h \sim_{(\text{dom}\ h') - U'} h'\ \wedge\ U' \cap (\text{dom}\ h) = \emptyset\ \wedge\ (n, r, U, h') \in mList)$
$\mathcal{F}\ fuse\ [x_0, x_1] \equiv \lambda E\ h\ h'\ v\ p.\ \forall n_0\ n_1.$
   $(\exists r_0\ U_0.\ E\ x_0\ =\ r_0\ \wedge\ (n_0, r_0, U_0, h)\ \in mList)\ \wedge$
   $(\exists\ r_1\ U_1.\ E\ x_1\ =\ r_1\ \wedge\ (n_1, r_1, U_1, h)\ \in\ mList)\ \wedge\ n_0 = n_1) \longrightarrow$
      $\mid \text{dom}\ h' \mid \leq \mid \text{dom}\ h \mid +n_0 + 1\ \wedge$
      $(\exists n'\ r'\ U'.\ v = r'\ \wedge\ (n', r', U', h') \in mList\ \wedge\ n' = n_0)$

In order to prove these two assertions, we first construct for each function a context, with specifications for all function calls in its body. We then prove the good context predicate, for these contexts. In both functions we find one direct recursive call, and therefore construct a one-element context for each function.

Lemma 1. *goodContext* $\mathcal{F}$ {(*copy* $\vec{xs}$, $\mathcal{F}$ *copy* $\vec{xs}$)} $\wedge$
   goodContext $\mathcal{F}$ {(*fuse* $\vec{xs}$, $\mathcal{F}$ *fuse* $\vec{xs}$)}

Proof structure. The same structure as in the proof for Theorem 3, but using VDMAx when encountering the recursive call. □

Now the resource consumption can be proven in an empty context:

Theorem 3. $\emptyset \rhd copy\ \vec{xs}\ :\ \mathcal{F}\ copy\ \vec{xs}\ \wedge\ \emptyset \rhd fuse\ \vec{xs}\ :\ \mathcal{F}\ fuse\ \vec{xs}$

Proof structure. In both clauses, by first applying the syntax-directed rules, and by applying CTXTWEAK, VDMADAPT and Lemma 1 when encountering the recursive call. The proof of the remaining subgoal proceeds by case distinction over the structure of the input, and over the cases in the conditional[4]. □

On coordination layer the main theorem below states that at each point in the execution any wire and any expression-layer result is either empty ($\bot$), or contains a reference to a list structure with N elements ($list \ldots$).

Theorem 4. $\square$
$$
\left(
\begin{array}{l}
(w1 = \bot \vee\ list\ h\ w1\ N)\ \wedge\ (w2 = \bot \vee\ list\ h\ w2\ N) \\
\wedge\ (w3 = \bot \vee\ list\ h\ w3\ N)\ \wedge\ (w4 = \bot \vee\ list\ h\ w4\ N) \\
\wedge\ (copy\_res1 = \bot \vee\ list\ h\ copy\_res1\ N) \\
\wedge\ (copy\_res2 = \bot \vee\ list\ h\ copy\_res2\ N) \\
\wedge\ (fuse\_res = \bot \vee\ list\ h\ fuse\_res\ N)
\end{array}
\right)
$$

Proof structure. Standard TLA reasoning reduces the proof of each conjunct as a separate invariant, which is outlined below[5]. □

We will focus on the proof of one conjunct, shown as Lemma 5. The other conjuncts are verified similarly. We start with the proof for an auxiliary theorem, relating the expression-layer result to the coordination-layer structure.

Lemma 2. $list\ H\ V\ M \wedge (W, H', S) = exe\ \emptyset(x := V)\ H\ (copy\ [x])\ \longrightarrow\ list\ H'\ V\ M$

Proof structure. This is proven from the expression-layer specification of *copy* using the Integration Theorem (vdmexe). □

Note that this property asserts that after box-execution the input V is still a *list*. Capital letters represent free variables. The invariant in Lemma 3 ensures that no other boxes (incl. the environment) can execute when $copy\_res1 \neq \bot$.

Lemma 3. $\square(copy\_res1 \neq \bot \longrightarrow w1 = \bot \wedge\ w2 = \bot \wedge w3 = \bot \wedge\ fuse\_res = \bot)$

Proof. The proof follows the 'standard Hume/TLA invariant proof structure' and follows from $\square(w3 \neq \bot \longrightarrow w1 = \bot \wedge\ w2 = \bot \wedge copy\_res1 = \bot \wedge copy\_res2 = \bot \wedge\ fuse\_res = \bot)$, which can be proven directly using the same structure. □

The following invariant asserts that wires and expression-layer results are well-formed, and that box executions do not interfere.

Lemma 4. $\square(w3 = \bot \vee\ list\ h\ w3\ N)$

Proof structure. This is proved using the 'standard Hume/TLA invariant proof structure' with the same strengthening as in the proof of Lemma 3. □

---

[4] These are theorems `copy_resources_ok` and `fuse_resources_ok` in `ListCopy.thy`.
[5] This is the main theorem in the theory file `ST.thy`.

Lemma 5. $\Box(copy\_res1 = \bot \lor \ list \ h \ copy\_res1 \ N)$

Proof. The proof follows the 'standard Hume/TLA invariant structure', strengthened by Lemmas 3 and 4, instantiating the free variables of Lemma 2. $\Box$

The proofs of the other conjuncts of Theorem 4 have the same structure, using lemmas over the expression-layer result of *copy* and *fuse* similar to Lemma 2. Since the entire expression-layer heap is discarded after execution, the size of the live heap at coordination-layer is the sum of all data structures reachable from the wires and results. From the 7 *list* predicates in Theorem 4 we can immediately conclude that $7N + 7$ is such a bound. Moreover, Lemma 3 and similar lemmas for the other expression-layer results restrict the possible combinations of non-$\bot$ values in the wires and results: after executing *copy* only $w1, w2, copy\_res1$ and $copy\_res2$ are non-$\bot$; after executing *fuse* only $w4$ and $fuse\_res$ are non-$\bot$. Therefore, we can refine the upper bound on the size of the live heap to $4N + 4$.

The (unoptimised) proof of *copy* takes 329 steps, with several instantiations of heap structures, address sets and list predicates. In particular, establishing the validity of the separation clause is the most complex part of the proof. To tackle this issue it would be possible to enhance the expression-level logic with concepts from separation logic [20], which enables reasoning about heap structures, without having to explicitly encode that the rest of the heap is unmodified.

6   Related Work

Our VDM-style, resource-aware program logic builds on our logic in [1], especially in handling mutual recursion and parameter adaptation, where the source language was an abstraction of JVM bytecode, but without algebraic data-types or higher-order functions as in Hume. In our formalisation we use techniques studied in Nipkow's Hoare-style logic for a simple while language in [18] and the Hoare-style logic for Java-light by von Oheimb [22], both formalised in Isabelle/HOL. Our choice of a VDM-style logic was partially driven by the study of variants of Hoare-style and VDM-style logics explored in Kleymann's thesis [13]. Isabelle/HOL formalisations of a program logic for a call-by-value, functional language are studied in [15]. Other related reasoning infrastructures, mostly for object-oriented languages, are: the Jive system [17], building on a Hoare-style logic for a Java subset [19]; the LOOP project's encoding of a Hoare-style logic, with extensions for reasoning about abrupt termination and side-effects, in PVS/Isabelle [11]; the *Why* system [4], with *Krakatoa* [16] as front-end, using Coq for formalisation and interactive proofs of JML-annotated Java programs.

Our choice of TLA for the coordination layer comes from the stateful nature of both systems, whereas approaches such as process algebras are stateless. We have previously used TLA to reason about scheduling algorithms [7]. In a shallow embedding of the expression layer directly in Isabelle/HOL, we have managed to automate the verification of 'standard Hume/TLA invariant proof structure' by developing a set of special purpose tactics. However, this does not enable the

type of resource reasoning achieved here, and we have no soundness guarantees of the embedding. Details can be found in the second author's PhD thesis [6]

Our integration of the two formalisations follows the "state-behaviour integration" approach, exemplified by csp2b [3], CSP||B [21] and CSP-OZ [5]. However, in this integration computation is state based (B or Z), and the communication uses a stateless *process algebra* [2]. In Hume on the other hand, computation is in a (stateless) purely functional language, and communication is in a (stateful) finite state machine language. All these tools are motivated by using suitable tools for the different aspects of a system. csp2b [3] works by translating the CSP (process algebra) [10] into a B machine, while CSP||B [21] and CSP-OZ [5] give a CSP semantics to the computational aspect (B and Object-Z). In our case, both the TLA and VDM representations are built on top of Isabelle/HOL, and integration is naturally achieved by "unlifting" into the Isabelle/HOL level, and no (resource preserving) translation is required.

7   Conclusion and Future Work

We have presented an integrated Isabelle/HOL formalisation of the two language layers of Hume. Our formalisation combines a high level of abstraction through the VDM-style program logic for the expression layer, ensured through the meta-theoretic soundness result, with a direct, shallow embedding of a TLA logic for the coordination layer. By formalising the most suitable style of logics for these two layers in a theorem prover, we obtain a powerful reasoning infrastructure. The expression layer soundness result has been formalised in the theorem prover. We have applied this infrastructure to one Hume program, verifying a basic resource property of the code, refining the latter through proofs of the availability of wire-values. This example shows that while the proofs can become lengthy, the proof strategies that can be applied follow naturally from the respective logics.

Notably, the TLA proofs follow a common structure, and we have explored their automation, through proof tactics, in the stand-alone coordination layer embedding [6]. In our integrated system, the proofs are (manually) driven by this structure at the coordination layer. Where needed, we make use of the VDM logic to prove the required properties for expressions. For resource properties only simple inequalities, e.g. over heap sizes, have to be proven. As discussed at the end of Section 5, the main structure of the proof follows naturally from the rules of the logics and the integration theorem. One complication, however, is to assert disjointness of data structures, requiring explicit reasoning over sets of addresses. Here techniques from separation logic could help [20] to raise the level of abstraction further. Finally, note that the Isabelle/HOL mechanisation currently only works for Isabelle 2007, and we are in the process of porting it to the latest version of Isabelle (Isabelle 2013), which we plan to publish in the Archive of Formal Proofs (http://afp.sourceforge.net). The theories for this paper are available at: http://www.macs.hw.ac.uk/~dsg/projects/hume-reasoning.html#FOPARA13

References

1. D. Aspinall, L. Beringer, M. Hofmann, H-W. Loidl, and A. Momigliano. A Program Logic for Resources. *Theoretical Computer Science*, 389(3):411–445, 2007.
2. J. C. M. Baeten. A Brief History of Process Algebra. *Theoretical Computer Science*, 335(2-3):131–146, 2005.
3. M.J. Butler. csp2B: A Practical Approach to Combining CSP and B. *Formal Aspect of Computing*, 12(3):182–198, 2000.
4. J.-C. Filliâtre. Why: a Multi-language Multi-prover Verification Tool. Research Report 1366, LRI, Université Paris Sud, March 2003.
5. C. Fischer. CSP-OZ: A combination of Object-Z and CSP. In *Formal Methods for Open Object-Based Distributed Systems (FMOODS '97)*, pages 423–438, 1997.
6. G. Grov. *Reasoning about Correctness Properties of a Coordination Programming Language.* PhD thesis, Heriot-Watt University, 2009.
7. G. Grov, G. Michaelson, and A. Ireland. Formal Verification of Concurrent Scheduling Strategies using TLA. In *International Conference on Parallel and Distributed Systems (ICPADS'07)*, pages 1–6, Hsinchu, Taiwan, December 2007. IEEE.
8. Gudmund Grov and Stephan Merz. A Definitional Encoding of TLA* in Isabelle/Hol. *Archive of Formal Proofs*, November 2011. http://afp.sf.net/entries/TLA, Formal proof development.
9. K. Hammond and G.J. Michaelson. Hume: a Domain-Specific Language for Real-Time Embedded Systems. In *Conf. Generative Programming and Component Engineering (GPCE '03)*, LNCS 2830, pages 37–56. Springer-Verlag, September 2003.
10. C. A. R. Hoare. *Communicating Sequential Processes*. Prentice-Hall, 1985.
11. M. Huisman and B. Jacobs. Java Program Verfication via a Hoare Logic with Abrupt Termination. In *Fundamental Approaches to Software Engineering (FASE'00)*, LNCS 1783, pages 284–303, 2000.
12. C. Jones. *Systematic Software Development Using VDM*. Prentice Hall, 1990.
13. T. Kleymann. *Hoare Logic and VDM: Machine-Checked Soundness and Completeness Proofs*. PhD thesis, LFCS, University of Edinburgh, 1999.
14. L. Lamport. The Temporal Logic of Actions. *ACM Transactions on Programming Languages and Systems*, 16(3):872–923, May 1994.
15. J. Longley and R. Pollack. Reasoning About CBV Functional Programs in Isabelle/HOL. In *Theorem Proving in Higher Order Logics (TPHOLs'04)*, LNCS 3223, pages 201–216. Springer, 2004.
16. C. Marché, C. Paulin-Mohring, and X. Urbain. The KRAKATOA Tool for Certification of JAVA/JAVACARD Programs Annotated in JML. *Journal of Logic and Algebraic Programming*, 58(1–2):89–106, 2004.
17. P. Müller, J. Meyer, and A. Poetzsch-Heffter. Programming and Interface Specification Language of JIVE. Technical report, Fernuniversität Hagen, 2000.
18. T. Nipkow. Hoare Logics for Recursive Procedures and Unbounded Nondeterminism. In *Computer Science Logic (CSL'02)*, LNCS 2471, pages 103–119, 2002.
19. A. Poetzsch-Heffter and P. Müller. A Programming Logic for Sequential Java. In *European Symp. on Programming (ESOP'99)*, LNCS 1576, pages 162–176, 1999.
20. J. Reynolds. Separation Logic: A Logic for Shared Mutable Data Structures. In *Symposium on Logic in Computer Science (LICS'02)*, pages 55–74, Copenhagen, July 2002. IEEE Computer Society.
21. S. Schneider and H. Treharne. CSP Theorems for Communicating B Machines. *Formal Aspects of Computing: Applicable Formal Methods*, 17(4):390–422, 2005.
22. D. von Oheimb. Hoare Logic for Java in Isabelle/HOL. *Concurrency and Computation: Practice and Experience*, 13(13):1173–1214, 2001.