

Contextual trace refinement for concurrent objects: Safety and progress

Brijesh Dongol¹ and Lindsay Groves²

¹ Brunel University London, UK
Brijesh.Dongol@brunel.ac.uk

² Victoria University of Wellington, New Zealand
lindsay@ecs.vuw.ac.nz

Abstract. Correctness of concurrent objects is defined in terms of safety properties such as linearizability, sequential consistency, and quiescent consistency, and progress properties such as wait-, lock-, and obstruction-freedom. These properties, however, only refer to the behaviour of the object in isolation, which does not tell us what guarantees these correctness conditions on concurrent objects provide to their client programs. This paper investigates the links between safety and progress properties of concurrent objects and a form of trace refinement for client programs, called contextual trace refinement. In particular, we show that linearizability together with a minimal notion of progress are sufficient properties of concurrent objects to ensure contextual trace refinement, but sequential consistency and quiescent consistency are both too weak. Our reasoning is carried out in the action systems framework with procedure calls, which we extend to cope with non-atomic operations.

1 Introduction

Concurrent objects provide operations that can be executed simultaneously by multiple threads, and provide a layer of abstraction to programmers by managing thread synchronisation on behalf of client programs, which in turn improves safety and efficiency. Correctness of concurrent objects is usually defined in terms of the possible *histories* of invocation and response events generated by executing the operations of a sequential specification object. There are several notions of safety for concurrent objects [12, 7]: sequential consistency, linearizability, and quiescent consistency being the most widely used. Similarly, there are many different notions of progress [12, 13], e.g., wait-, lock- and obstruction-freedom are popular non-blocking conditions.

Both safety and progress properties are stated in terms of a concurrent object in isolation, and disregard their context, i.e., the client programs that use them. Programmers (i.e., client developers) have therefore relied on informal “folk theorems” to link correctness conditions on concurrent objects and substitutability of objects within client programs. We seek to provide a formal account of this relationship, addressing the question: “Provided concurrent object OC is correct with respect to sequential object OA , how are the behaviours of $\mathcal{C}[OA]$ related to those of $\mathcal{C}[OC]$?”, where $\mathcal{C}[O]$ denotes a client program \mathcal{C} that uses object O , for different notions of correctness. One of the first

answers to this question was given by *abstraction theorems* [9], linking safety properties: sequential consistency and linearizability to a contextual notion of correctness called *observational refinement*, which defines substitutability with respect to the initial and final state of a system’s execution. For *terminating* clients, linearizability is shown to be equivalent to observational refinement, while sequential consistency is shown to be equivalent to observational refinement provided that clients do not communicate outside the given objects.

Since non-termination is common in many concurrent systems, e.g., operating systems and real-time controllers, our work aims to understand the link between concurrent correctness and substitutability for *potentially non-terminating* clients. Related to this aim is the work of Gotsman and Yang [10] and Liang et al. [15], who link observational refinement to safety and progress properties of concurrent objects. However, both [10] and [15] assume that the concurrent objects in question are already linearizable; in contrast, we do not assume linearizability. Further, [10] aims to understand *compositionality* of progress properties, while [15] develops *characterisations* of progress properties based on their observational guarantees.

The motivation for our work differs from [10, 15] in that we take *contextual trace refinement* as the underlying correctness condition when substituting *OC* for *OA* in \mathcal{C} , then aim to understand the safety/progress properties on *OC* that are required to guarantee trace refinement between $\mathcal{C}[OA]$ and $\mathcal{C}[OC]$. To this end, we develop an *action systems* framework that integrates and extends existing work [18, 1] from the literature, building on our preliminary results on this topic [8]. As part of our contributions we (i) extend Sere and Waldén’s treatment of action systems with procedures [18] with *non-atomic procedures*; (ii) develop a theory for *contextual trace refinement*, adapting Back and von Wright’s [1] theory for trace refinement of action systems, then reduce system-wide proof obligations (i.e., properties of the client and object together) to proof obligations on the objects only; (iii) show that linearizability [14] and minimal progress [13] together are sufficient to guarantee contextual trace refinement; and (iv) show that both sequential consistency and quiescent consistency are too weak to ensure contextual trace refinement, even in the presence of minimal progress.

2 Concurrent objects and their clients

We motivate concurrent objects using Treiber’s stack (Section 2.1). An example stack client (Section 2.2) is used to motivate contextual trace refinement (Section 2.3).

2.1 Client-object systems

We consider concurrent systems where a client consists of multiple threads which interact with one or more concurrent objects and shared variables. For example, the following client program consists of threads 1 and 2 using a shared stack *s*, and variables *x*, *y* and *z*.

```

Init x, y, z = 0, 0, 0
Thread 1:
  T1: s.push(1);
  T2: s.push(2);
  T3: s.pop(x);
Thread 2:
  U1: s.pop(y);
  U2: z := x;
```

```

Init: S = ⟨⟩
push(v) ==
atomic { S := ⟨v⟩^S }
pop ==
atomic {
  if S = ⟨⟩
  then return empty
else
  lv := head(S);
  S := tail(S);
  return lv }

```

Fig. 1. Abstract stack

```

Init: Head = null
push(v) ==
H1: n := new(Node);
H2: n.val := v;
    repeat
H3:  ss := Head;
H4:  n.next := ss;
H5: until
    CAS(Head,ss,n)
H6: return
pop ==
    repeat
P1: ss := Head;
P2: if ss = null
P3: then return empty
    else
P4:  ssn := ss.next;
P5:  lv := ss.val
P6: until
    CAS(Head,ss,ssn);
P7: return lv

```

Fig. 2. The Treiber stack

Thread 1 pushes 1 then 2 onto the stack s , then pops the top element of s and stores it in x . Concurrently, thread 2 pops the top element of s and stores it in y , then reads the value of x and stores it in z .

The abstract behaviour of a stack is defined in terms of a sequential object, as shown in Fig. 1. The abstract stack consists of a sequence of elements S together with two operations *push* and *pop* ($\langle \cdot \rangle$ and $\langle \cdot \rangle$ delimit sequences, $\langle \rangle$ denotes the empty sequence, and $\langle \cdot \rangle$ denotes sequence concatenation). Note that when the stack is empty, *pop* returns a special value *empty* that cannot be pushed onto the stack.

If concurrent objects are implemented using fine-grained concurrency, the call statements in their clients are not necessarily atomic because they may invoke non-atomic operations. Furthermore, depending on the implementation of s , we will get different traces of the client program because the effects of the concurrent operations on s may take effect in different orders. For example, Fig. 2 presents a simplified version of a non-blocking stack example due to Treiber [19]. In this implementation, each line of the *push* and *pop* corresponds to a single atomic step. Synchronisation of *push* and *pop* operations is achieved using a compare-and-swap (CAS) instruction, which takes as input a (*shared*) variable gv , an *expected value* lv and a *new value* nv :

```

CAS(gv, lv, nv) ≐ atomic { if (gv = lv)
                          then gv := nv ; return true
                          else return false }

```

With this stack implementation, the executions of operations, say T1 and U1, in the above client may overlap, and different behaviours may be observed according to the order in which steps of the different threads are executed. Treiber's stack is linearizable with respect to the abstract stack in Fig. 1, so the effect of each operation call takes place between its invocation and its response. If a different stack implementation is used which satisfies a more permissive correctness condition, such as sequential consistency or quiescent consistency [12], a wider range of behaviours may be observed by its client.

2.2 Observability and contextual trace refinement

With an example client-object system in place, we return to the main question for this paper: What guarantees do correctness conditions on concurrent objects provide

to clients that use the objects? Furthermore, how can one address divergence, termination and reactivity of a client? To address these, we first pin down the aspects of the system being developed that are visible to an external observer. Following Filipović et al. [9], we take the state of the client variables to be observable, and the state of the objects they use to be unobservable. Therefore, for the client program in Section 2.1, variables x , y and z are observable, but none of the variables of the stack implementation s are observable. This allows us to reason about a client with respect to different implementations of s . Second, we define *when* a system may be observed. Unlike Filipović et al. [9] who only observe the client state at the beginning and end of a client’s execution, we assume that the states *throughout* a client’s execution are visible. This allows us to accommodate, for example, reactive clients, which interact with an observer in some way even if they are potentially non-terminating.

Therefore, our notion of correctness for the combined client-object system will be a form of *observational refinement* that holds iff every (observable) trace of a client using a concurrent object is equivalent to some (observable) trace of the same client using the corresponding abstract specification of the object. The end result is that from the perspective of a client program, it will be impossible to tell whether it is using the concurrent object, or its abstract (sequential) specification.

Example 1. Let \mathcal{D} denote the client program in Section 2.1, TS denote the Treiber stack in Fig. 2, and AS denote the abstract stack in Fig. 1. Suppose the stack s in \mathcal{D} is an instance of TS . Then the following is a possible observable trace of $\mathcal{D}[TS]$:

$$tr \hat{=} \langle (x, y, z) \mapsto (0, 0, 0), (x, y, z) \mapsto (0, 2, 0), (x, y, z) \mapsto (1, 2, 0), (x, y, z) \mapsto (1, 2, 1) \rangle$$

where $(x, y, z) \mapsto (0, 0, 0)$ is shorthand for the state $\{x \mapsto 0, y \mapsto 0, z \mapsto 0\}$, and we ignore *stuttering*, i.e., consecutive states that leave the observable state unchanged. Trace tr is obtained by initialising as specified by `Init`, then executing `T1`, `T2`, `U1`, `T3`, then `U2` to completion; i.e. they execute their operation call without interruption. It is straightforward to see that tr can also be generated by $\mathcal{D}[AS]$, i.e., when using the abstract stack for s . Thus tr can be accepted as being correct. Executions can, of course, be much more complicated than tr — because TS consists of non-atomic operations, executions of `T1`, `T2` or `T3` may overlap with `U1` or `U2`. \square

We say that TS *contextually trace refines* AS with respect to the client program \mathcal{C} iff every trace of $\mathcal{C}[TS]$ is a possible trace of $\mathcal{C}[AS]$. In this paper, we wish to know whether contextual refinement holds for every client program. To this end, we say TS *contextually trace refines* AS iff TS contextually trace refines AS with respect to every client program \mathcal{C} .

2.3 Correctness conditions on concurrent objects

There are many notions of correctness for concurrent objects, and these are defined in terms of *histories* of invocation and response events, corresponding to operation calls on the object [12] (see Section 5 for details).

Concurrent histories may consist of both overlapping and non-overlapping operation calls, inducing a partial order on events. Safety properties define how, if at all, this

partial order is preserved by the corresponding abstract histories generated by the corresponding sequential object [12, 7]. We will consider three different safety properties. *Sequential consistency* is a simple condition requiring the order of operation calls in a concrete history for a single process to be preserved. Operation calls performed by different processes may be reordered in the abstract history even if the operation calls do not overlap in the concrete history. *Linearizability* strengthens sequential consistency by requiring the order of non-overlapping operations to be preserved. Operation calls that overlap in the concrete history may be reordered when mapping to an abstract history. *Quiescent consistency* is weaker than linearizability, but is incomparable to sequential consistency. A concurrent object is said to be quiescent at some point in its history if none of its operations are executing at that point. Quiescent consistency requires the order of operation calls that are separated by a quiescent point to be preserved. Operation calls that are not separated by a quiescent point may be reordered, including operations performed by the same process.

Progress conditions on concurrent objects are necessary to ensure that clients will eventually be able to continue execution after calling operations on the objects they use. We consider a notion of progress called *minimal progress* [13], which guarantees that after some finite number of steps, some operation of the concurrent object terminates.

3 Modelling client-object systems

Our formal framework for reasoning about contextual trace refinement is based on existing work on action systems with procedures [18], which we extend to cope with potentially non-atomic operations. We let Var and Val denote the types of variables and values, respectively. We distinguish between *unobservable* and *observable* variables using Var_U and Var_O , respectively, where $Var_U, Var_O \subseteq Var$ and $Var_U \cap Var_O = \emptyset$. A *state* is a function $\Sigma_V \hat{=} V \rightarrow Val$, where $V \subseteq Var$, and a *predicate* of type K is of type $\mathcal{P}K \hat{=} K \rightarrow \mathbb{B}$, e.g., a *state predicate* over V is of type $\mathcal{P}\Sigma_V$.

The abstract syntax of an action system is of the form:

$$\mathcal{A} ::= \llbracket \mathbf{var}_u L; \mathbf{var}_o G; \mathbf{proc} \, ph_1 = P_1 \dots \mathbf{proc} \, ph_n = P_n; I; \mathbf{do} \, A \, \mathbf{od} \rrbracket$$

where $L \subseteq Var_U$ is a set of *unobservable variables* and $G \subseteq Var_O$ a set of *observable variables*; each $ph_i = P_i$ is a (non-recursive) procedure declaration; I is an action modelling initialisation; and A is the main action. Within each $ph_i = P_i$, P_i is an action and ph_i is a procedure heading $p_i(\mathbf{val} \, v, \mathbf{res} \, x)$ with procedure name p_i and optional call-by-value and call-by-result parameters v and x . Procedure declarations may additionally be parameterised by thread identifiers.

The abstract syntax of *actions* is of the form:

$$A ::= \mathbf{var} \, x \mid \mathbf{rav} \, x \mid \mathbf{skip} \mid x \in E \mid x := e \mid p(in, out) \mid A_1; A_2 \mid b \rightarrow A \mid A_1 \sqcap A_2$$

where x is a variable, E is a set-valued expression, e is an expression, p is a procedure name, in and out are inputs and outputs to a procedure (which may be a value or a variable), and b is a predicate. Actions $\mathbf{var} \, x$ and $\mathbf{rav} \, x$ introduce and remove variable x from the state space, respectively, \mathbf{skip} is an action that leaves the state unchanged, $x \in E$ denotes non-deterministic assignment, $x := e$ denotes assignment, $p(e, x)$ is a procedure call with value parameter e and result parameter x , $A_1; A_2$ is sequential composition of

A_1 and A_2 , $b \rightarrow A$ is a guarded action, and $A_1 \sqcap A_2$ is (demonic) choice between A_1 and A_2 .

The meaning of *parameterless procedures* is given by syntactically replacing each procedure call p in A by the procedure body, P . Procedure parameters are handled by introducing new local variables with the same name; for call-by-value, the new variable is initialised with the value of the actual parameter, while for call-by-results, the final value is copied to the variable passed as the parameter (see [18]). We give examples of these in Examples 2 and 3 below.

When invoking non-atomic operations, it will be important to detect when the invoked operation has terminated. To this end, we assume that a variable \widehat{pc}_t is used to control the flow of execution within an operation; thus \widehat{pc}_t must be declared whenever thread t is currently executing an operation. Formally, we use state predicate

$$dec.v \hat{=} \lambda \sigma \bullet v \in dom(\sigma)$$

which holds iff variable v is declared in the domain of the given state. We use ‘.’ for function application.

Example 2. Consider again the client program \mathcal{D} from Section 2.1 and suppose it uses the abstract stack object AS in Fig. 1. The action system modelling the client-object system is $\mathcal{D}[AS]$, given below. The shared stack is a sequence modelled by an unobservable variable S . The client consists of variables x, y and z , as well as program counters pc_1 and pc_2 (which we distinguish from \widehat{pc}_t). We assume

$$npc_t(k) \hat{=} (dec.\widehat{pc}_t \rightarrow \mathbf{skip}) \sqcap (\neg dec.\widehat{pc}_t \rightarrow pc_t := k)$$

is an action that sets pc_t to k if t completes the operation it is currently executing.

[[**var**_u S ; **var**_o x, y, z, pc_1, pc_2 ;
proc $push_t(\mathbf{val} \textit{in}) = S := \langle \textit{in} \rangle \wedge S$
proc $pop_t(\mathbf{res} \textit{out}) = S = \langle \rangle \wedge \neg dec.\widehat{pc}_t \rightarrow \mathbf{var} \textit{ret}, \widehat{pc}_t; \textit{ret} := \textit{empty}; \widehat{pc}_t := 1$
 $\sqcap S \neq \langle \rangle \wedge \neg dec.\widehat{pc}_t \rightarrow \mathbf{var} \textit{ret}, \widehat{pc}_t;$
 $\textit{ret}, S := \textit{head}.S, \textit{tail}.S; \widehat{pc}_t := 1$
 $\sqcap \widehat{pc}_t = 1 \rightarrow \textit{out} := \textit{ret}; \mathbf{rav} \textit{ret}, \widehat{pc}_t;$
 $S, pc_1, pc_2 := \langle \rangle, T1, U1; x, y, z := 0, 0, 0;$
do $pc_1 = T1 \rightarrow push_1(1); npc_1(T2)$
 $\sqcap pc_1 = T2 \rightarrow push_1(2); npc_1(T3)$
 $\sqcap pc_1 = T3 \rightarrow pop_1(x); npc_1(\perp)$
 $\sqcap pc_2 = U1 \rightarrow pop_2(y); npc_2(U2)$
 $\sqcap pc_2 = U2 \rightarrow z, pc_2 := x, \perp$ **od**]]

□

Example 3. The $push_t$ operation of the Treiber stack is defined as follows. We assume $newNode.n \hat{=} n \notin Nodes$; $Nodes := Nodes \setminus \{n\}$ assigns n to be a new node from the available set of nodes $Nodes$. For simplicity, we assume $Nodes$ is an infinite set (e.g., the natural numbers), so a new node is always available. Thus we have:

proc $push_t(\mathbf{val} \textit{in}) = \neg dec.\widehat{pc}_t \rightarrow \mathbf{var} \widehat{pc}_t, v_t, n_t, ss_t; v_t := \textit{in}; \widehat{pc}_t := H1$
 $\sqcap \widehat{pc}_t = H1 \rightarrow newNode.n_t; \widehat{pc}_t := H2$
 \dots
 $\sqcap \widehat{pc}_t = H6 \rightarrow \mathbf{rav} \widehat{pc}_t, v_t, n_t, ss_t$

The *pop* operation is similar, except that it additionally sets the output variable to the returned value.

$$\begin{aligned} \mathbf{proc} \text{ pop}_t(\mathbf{res} \text{ out}) &= \neg \text{dec}.\widehat{pc}_t \rightarrow \mathbf{var} \widehat{pc}_t, ss_t, ssn_t, lv_t; \widehat{pc}_t := P1 \\ &\quad \dots \\ &\quad \sqcap \widehat{pc}_t = P7 \rightarrow \text{out} := lv_t; \mathbf{rav} \widehat{pc}_t, ss_t, ssn_t, lv_t \end{aligned}$$

The action system resulting from using the Treiber stack (which we will refer to as *TS*) as the shared concurrent object in Section 2.1 is $\mathcal{D}[TS]$. It is similar to the action system in Example 2, except that the unobservable variables are *Nodes* (the set of all available nodes), *Head* (a pointer to a node, or *null*), *val* (a partial function of type $Nodes \rightarrow Val$), *next* (a partial function of type $Nodes \rightarrow Node$); the procedure declarations above are used; and initialisation of the object is $Nodes, Head, val, next := \mathbb{N}, null, \emptyset, \emptyset$. \square

We now make the concept of an object and the notation $\mathcal{C}[O]$ for an object O and client \mathcal{C} more precise. An *object* is a triple $O \hat{=} (L, P, I)$, where L is a set of variables, $P \hat{=} \{ph_{1,t} = P_{1,t}, \dots, ph_{n,t} = P_{n,t}\}$ is a set of (potentially parameterised) procedure declarations, and I is an initialisation action. A *client* is a triple $\mathcal{C} \hat{=} (G, A, J)$, where G is a set of variables, and A and J are the main and initialisation actions, respectively. Then $\mathcal{C}[O]$ is the action system

$$\llbracket [\mathbf{var}_u L; \mathbf{var}_o G; \mathbf{proc} \text{ } ph_{1,t} = P_{1,t} \dots \mathbf{proc} \text{ } ph_{n,t} = P_{n,t}; I; J; \mathbf{do} A \mathbf{od}] \rrbracket.$$

The next section formalises the semantics of action systems and defines our notion of contextual trace refinement for it.

4 Semantics and contextual trace refinement

We now give the semantics for action systems and define contextual trace refinement, which extends the existing theory on trace refinement [1]. Note that we only use part of the action systems framework. In particular, to develop a more direct link to trace refinement, we only give a relational semantics for actions.

We assume that expressions are functions from states to values. A *relation* is of type $\mathcal{R}(K, K') \hat{=} K \rightarrow \mathcal{P}K'$, thus a *state relation* is of type $\mathcal{R}(\Sigma_V, \Sigma_{V'})$, where $V, V' \subseteq \text{Var}$. Assume r, r_1 and r_2 are state relations, b is a predicate and S is a set. We let

- $(r_1 \circ r_2).\gamma.\gamma' \hat{=} \exists \gamma'' \bullet r_1.\gamma.\gamma'' \wedge r_2.\gamma''.\gamma'$ denote *relational composition*,
- $(b \triangleleft r).\gamma.\gamma' \hat{=} b.\gamma \wedge r.\gamma.\gamma'$ denote *domain restriction*, and
- $S \triangleleft r = \{(\gamma, \gamma') \in r \mid \gamma \notin S\}$ denote *domain anti-restriction*.

For a function f , we let $f \oplus \{x \mapsto v\} \hat{=} \lambda z \in \text{dom}(f) \bullet \mathbf{if} \ z = x \ \mathbf{then} \ v \ \mathbf{else} \ f.z$ denote *functional overriding*.

Definition 1. *The (relational) semantics of an action A is given by $\text{rel}.A$:*

$$\begin{aligned} \text{rel}.\mathbf{var} \ x &\hat{=} \lambda \sigma \bullet \lambda \sigma' \bullet & \text{rel}.\mathbf{skip} &\hat{=} \text{id} \\ &(\{x\} \triangleleft \sigma') = \sigma \wedge \text{dec}.x.\sigma' & \text{rel}.(b \rightarrow A_1) &\hat{=} b \triangleleft \text{rel}.A_1 \\ \text{rel}.\mathbf{rav} \ x &\hat{=} \lambda \sigma \bullet \lambda \sigma' \bullet (\{x\} \triangleleft \sigma) = \sigma' & \text{rel}.(A_1; A_2) &\hat{=} \text{rel}.A_1 \circ \text{rel}.A_2 \\ \text{rel}.(x := e) &\hat{=} \lambda \sigma \bullet \lambda \sigma' \bullet \sigma' = \sigma \oplus \{x \mapsto e.\sigma\} & \text{rel}.(A_1 \sqcap A_2) &\hat{=} \text{rel}.A_1 \vee \text{rel}.A_2 \\ \text{rel}.(x \in E) &\hat{=} \lambda \sigma \bullet \lambda \sigma' \bullet & & \\ &\exists k : E.\sigma \bullet \sigma' = \sigma \oplus \{x \mapsto k\} & & \end{aligned}$$

Recall that the semantics of a procedure call is given by substitution as described in Section 3. We let $grd.A.\gamma \hat{=} \gamma \in dom(rel.A)$ denote the *guard* of A . Because an action system is a loop with a non-deterministic choice over actions [1], we frequently use iteration in our reasoning. Formally, finite iteration of relation r (denoted r^*) is defined as follows:

$$r^0 \hat{=} id \quad r^{k+1} \hat{=} r \circ r^k \quad r^* \hat{=} \exists k \in \mathbb{N} \bullet r^k$$

The semantics of an iterated action is defined by lifting from iteration defined on relations, namely, $rel.A^* \hat{=} (rel.A)^*$. We say an *iterated execution of A terminates from state γ* iff $term.A.\gamma \hat{=} \exists k \bullet \forall \gamma' \bullet (rel.A)^k.\gamma.\gamma' \Rightarrow \neg grd.A.\gamma'$. Note that $\neg grd.A.\gamma \Rightarrow term.A.\gamma$ holds for all actions A and states γ .

We use $seq X$ to denote (possibly infinite) sequences of elements of type X , and assume indices start from 0.

Definition 2. A possibly infinite sequence of states s is a trace of action system \mathcal{A} iff $\exists \sigma \bullet rel.I.\sigma.(s.0) \wedge \forall i : dom(s) \setminus \{0\} \bullet rel.A.(s.(i-1)).(s.i)$ holds.

A trace is *complete* iff either the trace is of infinite length or the guard of A does not hold in the last state of the trace. The set of all *complete traces* of an action system \mathcal{A} is denoted $\llbracket \mathcal{A} \rrbracket$.

Traces (Definition 2) provide a conceptually simple model for a system's execution, and trace refinement provides a conceptually simple notion of substitutability [1]. Typically, because a concrete system is more fine-grained than the abstract, one must remove stuttering from a trace. An action system may also exhibit *infinite stuttering* by generating a trace that ends with an infinite sequence of consecutive stuttering steps. After infinite stuttering, one will never be able to observe any state changes, and hence, we treat infinite stuttering as *divergence*, which is denoted by a special symbol ' $\uparrow \notin \Sigma$ '. For any trace $s \in \llbracket \mathcal{A} \rrbracket$, we define $Tr.s$ to be the non-stuttering observable sequence of states, possibly followed by \uparrow , which is obtained from s as follows. First, we obtain a sequence s' by removing all finite stuttering in s and replacing any infinite stuttering in s by \uparrow . Second, for each $i \in dom(s')$, we let $(Tr.s').i = \mathbf{if } s'.i \neq \uparrow \mathbf{then } Var_U \triangleleft s'.i \mathbf{else } \uparrow$. It is straightforward to define functions that formalise both the steps above (see for example [6]).

Definition 3. Abstract action system \mathcal{A} is trace refined by concrete action system \mathcal{C} (denoted $\mathcal{A} \sqsubseteq \mathcal{C}$) iff $\forall s' \in \llbracket \mathcal{C} \rrbracket \bullet \exists s \in \llbracket \mathcal{A} \rrbracket \bullet Tr.s = Tr.s'$ holds.

Back and von Wright have developed simulation rules (details elided due to lack of space) for verifying trace refinement of action systems [1], which we adapt to reason about client-object systems in Lemmas 1 and 2. First, we formalise the meaning of contextual trace refinement. The notion is similar to the notion of data refinement given by He et al. [11, 3], but extended to traces, which enables one to cope with non-terminating reactive systems.

Definition 4. An abstract object OA is contextually trace refined by a concrete object OC , denoted $OA \sqsubseteq OC$, iff for any client C we have $C[OA] \sqsubseteq C[OC]$.

In this paper, for simplicity, we assume that (atomic) actions do not abort [3], therefore the proof obligations for aborting actions do not appear in Lemmas 1 and 2 below –

it is straightforward to extend our results to take aborting behaviour into account. However, like Back and von Wright [1], our notion of refinement ensures *total correctness* of the systems we develop, i.e., the concrete system may only deadlock (or diverge) if the abstract system deadlocks (or diverges). Thus, in addition to the standard step correspondence proof obligations for ensuring safety of the concrete system, we include Back and von Wright’s proof obligations that ensure progress.

Because the entire state of the client is observable, the proof obligations pertaining to the client can be trivially discharged, leaving one with proof obligations that only refer to the object. For procedure declarations $P \hat{=} \{ph_{1,t} = P_{1,t}, \dots, ph_{n,t} = P_{n,t}\}$, we let $tact.v.x.t.P \hat{=} p_{1,t}(v,x) \sqcap \dots \sqcap p_{n,t}(v,x)$ denote the choice between procedures in P for inputs v and x and thread t then define:

$$act.P \hat{=} \prod_{v,x,t} tact.v.x.t.P \quad rem.P \hat{=} \prod_{v,x,t} dec.\widehat{pc}_t \rightarrow tact.v.x.t.P$$

To simplify the syntax, we implicitly assume that in $tact.v.x.t.P$ the inputs v and x are of the correct type for each procedure. Guard $dec.\widehat{pc}_t$ is used to detect whether the procedure being executed by thread t has terminated — if t is executing a procedure, say $ph_{i,t}$, we know $dec.\widehat{pc}_t$ will hold and when this procedure terminates $\neg dec.\widehat{pc}_t$ will hold, which disables thread t . The intention is to use $rem.P$ in (4) below, which attempts to execute the remaining steps of the running operations by each thread to completion.

Lemma 1 (Forward simulation). *If $OA = (L_A, P_A, I_A)$ and $OC = (L_C, P_C, I_C)$ are objects, then $OA \sqsubseteq OC$ if there exists a relation R and the following hold for any states σ, τ and τ' :*

$$rel.I_C.\tau' \Rightarrow \exists \sigma \bullet R.\sigma.\tau' \wedge rel.I_A.\sigma \quad (1)$$

$$R.\sigma.\tau \wedge rel.(act.P_C).\tau.\tau' \Rightarrow \exists \sigma' \bullet R.\sigma'.\tau' \wedge rel.(act.P_A)^*.\sigma.\sigma' \quad (2)$$

$$R.\sigma.\tau \wedge \neg grd.(act.P_C).\tau \Rightarrow \neg grd.(act.P_A).\sigma \quad (3)$$

$$true \Rightarrow term.(rem.P_C).\tau \quad (4)$$

The first three proof obligations are straightforward. Proof obligation (4) requires that the main action of the concrete object OC terminates if threads do not invoke new operations after the operation currently being executed has terminated. Note that (4) does not rule out infinite stuttering within the program $\mathcal{C}[OC]$, but it does ensure that any infinite stuttering is caused by the client as opposed to the object OC , and hence, this infinite stuttering must also be present within $\mathcal{C}[OA]$. Therefore, if (4) holds, so does Back and von Wright’s non-termination condition.

Dually to forward simulation, there exists a method of *backward simulation*, which requires that the abstract action system under consideration is *continuous*. An action system \mathcal{A} with main action A is *continuous* iff for all σ , the set $\{\sigma' \mid rel.A.\sigma.\sigma'\}$ is finite, i.e., A does not exhibit infinite non-determinism.

Lemma 2 (Backward simulation). *Suppose $OA = (L_A, P_A, I_A)$ and $OC = (L_C, P_C, I_C)$ are objects and \mathcal{C} is a client such that $\mathcal{C}[OA]$ is continuous. Then $\mathcal{C}[OA] \sqsubseteq \mathcal{C}[OC]$ holds if there exists a total relation R and for any states σ' and τ, τ' condition (4) as well as each of the following hold:*

$$rel.I_C.\tau' \wedge R.\sigma'.\tau' \Rightarrow rel.I_A.\sigma' \quad (5)$$

$$rel.(act.P_C).\tau.\tau' \wedge R.\sigma'.\tau' \Rightarrow \exists \sigma \bullet R.\sigma.\tau \wedge rel.(act.P_A)^*.\sigma.\sigma' \quad (6)$$

$$\neg grd.(act.P_C).\tau \Rightarrow \exists \sigma \bullet R.\sigma.\tau \wedge \neg grd.(act.P_A).\sigma \quad (7)$$

Lemmas 1 and 2 reduce the proof obligations for trace refinement of client-object systems to the level of objects only. This allows one to explore properties of objects in isolation to guarantee contextual trace refinement.

5 Events and histories

This section provides background for defining safety (e.g., linearizability) and progress (e.g., lock-freedom) properties of concurrent objects [12]. We define both types of properties in terms of *histories* of invocation and response events [14, 12] that record the externally visible interaction between a client and the object it uses. The type of an event is *Event*, which is defined as follows [4]:

$$Event ::= inv\langle\mathbb{N} \times Op \times (Val \cup \{\perp\})\rangle \mid ret\langle\mathbb{N} \times Op \times (Val \cup \{\perp\})\rangle$$

The components of each event are the thread identifier, the operation name and input/output values. We use $\perp \notin Val$ to denote an invocation (return) event that has no input (output). Thus, for example, $inv(1, push, 2)$ denotes an *push* invocation by thread 1 with value 2, and $ret(1, push, \perp)$ denotes a return from this invocation.

The history of an object is a (potentially infinite) sequence of events, i.e., *History* $\hat{=}$ seq *Event*. A history of an object is generated by an execution of a *most-general client* in Definition 5 below, but first we describe how invocations and responses are recorded in a history. For an object $O \hat{=} (L, \{ph_{1,t} = P_{1,t}, \dots, ph_{n,t} = P_{n,t}\}, I)$ assuming $H \notin L$ is a history variable, we let $P_{i,t}^H$ be the *history-extended* procedure derived from $P_{i,t}$ by additionally recording invocation and response events in H (also see [4]).

Example 4. The history-extended procedure for $push_t$ from Example 2 is:

$$H := H \wedge \langle inv(t, push, in) \rangle; S := \langle in \rangle \wedge S; H := H \wedge \langle ret(t, push, \perp) \rangle$$

while the history-extended version of $push_t$ procedure from Example 3 is:

$$\neg dec.\hat{p}c_t \rightarrow \mathbf{var} \hat{p}c_t, v_t, n_t, ss_t; v_t := in;$$

$$H := H \wedge \langle inv(t, push, in) \rangle; \hat{p}c_t := H1$$

...

$$\sqcap \hat{p}c_t = H6 \rightarrow H := H \wedge \langle ret(t, push, \perp) \rangle; \mathbf{rav} \hat{p}c_t, v_t, n_t, ss_t \quad \square$$

Definition 5. The most general client of $O \hat{=} (L, \{ph_{1,t} = P_{1,t}, \dots, ph_{n,t} = P_{n,t}\}, I)$ is the action system $\mathcal{M}[O]$ below, where $H \notin L$ is its history, $tt \notin L$ is a fresh variable that models termination and $P^H \hat{=} \{ph_{1,t} = P_{1,t}^H \dots ph_{n,t} = P_{n,t}^H\}$ is the set of history extended procedures:

$$\mathcal{M}[O] \hat{=} \llbracket \mathbf{var}_u L \cup \{H, tt\}; \mathbf{var}_o Var_O;$$

$$\mathbf{proc} ph_{1,t} = P_{1,t}^H \dots \mathbf{proc} ph_{n,t} = P_{n,t}^H ;$$

$$I; H := \langle \rangle; tt := false ;$$

$$\mathbf{do} \neg tt \rightarrow act.P^H \sqcap (\prod_{w:V_O, a:Val} w := a) \sqcap tt := true \mathbf{od} \rrbracket$$

Thus, $\mathcal{M}[O]$ includes unobservable variables H (initially $\langle \rangle$) and tt (initially *false*), which model the history and termination of $\mathcal{M}[O]$, respectively. Provided tt is false, at each iteration of the action system either

- a step of a history-extended procedure of O is executed, or
- some observable variable is set to a non-deterministically chosen value, or
- $\mathcal{M}[O]$ terminates by setting tt to *true*.

The intention of $\mathcal{M}[O]$ is to model all possible client behaviours, including for instance faults (where a thread stops running) or a divergence (where a thread repeatedly executes the same operation).

Definition 6. *The set of histories of an object O is given by*

$$\{h \in \text{seq Event} \mid \exists s : \llbracket \mathcal{M}[O] \rrbracket \bullet \exists i : \text{dom}(s) \bullet h = (s.i).H\}.$$

6 Contextual trace refinement: Progress

The progress condition we will consider is *minimal progress*, which guarantees system-wide progress, even though there may be individual threads that may not make progress [13]. To formalise minimal progress, we say event e_1 *matches* e_2 iff $\text{matches}(e_1, e_2) \hat{=} \exists t, o, u, v \bullet e_1 = \text{inv}(t, o, u) \wedge e_2 = \text{ret}(t, o, v)$ holds, i.e., e_1 is an invocation of an operation by a thread and e_2 is the corresponding return. We say $m \in \text{dom}(h)$ is a *pending invocation* iff $\text{pi}(m, h) \hat{=} \forall n \in \text{dom}(h) \bullet m < n \Rightarrow \neg \text{matches}(h.m, h.n)$ holds.

An object O satisfies minimal progress iff for every trace tr of the $\mathcal{M}[O]$, it is always the case that in the future, either $\mathcal{M}[O]$ terminates, or there is some pending operation invocation that completes and returns.

Definition 7. *An object O satisfies minimal progress iff for every $s \in \llbracket \mathcal{M}[O] \rrbracket$ and $i \in \text{dom}(s)$, there exists a $j \in \text{dom}(s)$ such that $i \leq j$ and*

$$(s.j).tt \vee \exists m \bullet \text{pi}(m, (s.j).H) \wedge \neg \text{pi}(m, (s.(j+1)).H).$$

That is, for any trace s of $\mathcal{M}[O]$ and index $i \in \text{dom}(s)$ there is a state $s.j$ (where $j \geq i$) from which some pending operation in $s.j$ completes. There are a variety of objects that satisfy minimal progress, e.g., wait-, lock-free objects under any scheduler, and obstruction-free objects under isolating schedulers (see [13] for details). Objects that do not satisfy minimal progress include obstruction free implementations that are executed using a weakly fair scheduler.

The lemma below states that any object that satisfies minimal progress does not suffer from deadlock, and is guaranteed to terminate if no additional operations are invoked.

Lemma 3. *If $O = (L, P, I)$ satisfies minimal progress, then for any $\gamma \in \llbracket \mathcal{M}[O] \rrbracket$ and $i \in \text{dom}(\gamma)$, both $\text{grd}(\text{act}.P).(\gamma.i)$ and condition (4) hold.*

Using Lemma 3, we simplify and combine Lemmas 1 and 2. In particular, we are left with the proof obligations for safety only as in the theorem below.

Theorem 1. Suppose $OA = (L_A, P_A, I_A)$ and $OC = (L_C, P_C, I_C)$ are objects, OC satisfies minimal progress, and $R \in \mathcal{R}(\Sigma_{L_A}, \Sigma_{L_C})$. Then

1. $OA \hat{=} OC$ if both (1) and (2) hold, and
2. for any client \mathcal{C} such that $\mathcal{C}[OA]$ is continuous, $\mathcal{C}[OA] \sqsubseteq \mathcal{C}[OC]$ holds if R is total and both (5) and (6) hold.

7 Safety and contextual trace refinement

We give the formal definition of safety properties using the nomenclature in [7] and [4]. We say $m, n \in \text{dom}(h)$ form a *matching pair* in h iff $mp(m, n, h)$ holds, where $mp(m, n, h) \hat{=} m < n \wedge \text{matches}(h.m, h.n) \wedge \forall i \bullet m < i < n \Rightarrow \pi_1.(h.i) \neq \pi_1.(h.m)$ and π_i is the *projection function* returning the i th element of the given tuple.

Following [7], safety properties are defined in terms of a history h and a *mapping function* f between indices. The *sequential history* corresponding to h and f is obtained using $\text{map}(h, f) \hat{=} \{f(k) \mapsto h(k) \mid k \in \text{dom}(f)\}$. Different safety properties are defined by placing different types of restrictions on f . The most basic restriction is validity of a mapping. We say a function f is a *valid mapping function* if, for any history h , (a) the domain of f is contained in the domain of h , (b) the range of f is a consecutive sequence starting from 0, (c) f only maps matching pairs in h , and (d) matching pairs in h are mapped to consecutive events in the target abstract history. Assuming $[m, n]$ is the set of integers from m to n inclusive, we formalise validity for mapping functions using $\text{VMF}(h, f)$, where

$$\begin{aligned} \text{VMF}(h, f) \hat{=} & \text{dom}(f) \subseteq \text{dom}(h) \wedge (\exists n : \mathbb{N} \bullet \text{ran}(f) = [0, n - 1]) \wedge \text{injective}(f) \wedge \\ & (\forall m, n : \text{dom}(h) \bullet mp(m, n, h) \Rightarrow (m \in \text{dom}(f) \Leftrightarrow n \in \text{dom}(f))) \wedge \\ & (\forall m, n : \text{dom}(f) \bullet mp(m, n, h) \Rightarrow f.n = f.m + 1) \end{aligned}$$

When formalising correctness conditions, one must also consider *incomplete histories*, which have pending operation invocations that may or may not have taken effect. To cope with these, like Herlihy and Wing [14], we use *history extensions*, which are constructed from a history h by concatenating a sequence of returns corresponding to some of the pending invocations of h . A *correctness condition* Z is a predicate on a history and a mapping function.

Definition 8. A concurrent object OC implementing an abstract object OA is correct with respect to a correctness condition Z , denoted $OC \models_{OA} Z$, iff for any history h of OC , there exists an extension he of h , a valid mapping function f such that $\text{VMF}(he, f) \wedge Z(he, f)$ holds and $\text{map}(he, f)$ is a history of OA .

7.1 Linearizability

We now show that linearizability is a sufficient safety condition for discharging the proof obligations in Theorem 1. Linearizability is a *total* condition, which means that all completed (i.e., returned) operation calls in a given history h must be mapped by f .³ In addition, it must satisfy an *order* condition *lin*, which states that the return of an

³ This is in contrast to *partial* conditions defined for relaxed memory (see [7] for details).

operation may not be reordered with an invocation that occurs after it. We use $inv?(e) \triangleq \exists t, o, v \bullet e = inv(t, o, v)$ if e is an invocation event and $ret?(e) \triangleq \exists t, o, v \bullet e = ret(t, o, v)$ if e is a response.

$$\begin{aligned} total(h, f) &\triangleq \forall m : dom(h) \bullet \neg pi(m, h) \Rightarrow m \in dom(f) \\ lin(h, f) &\triangleq \forall m, n : dom(f) \bullet m < n \wedge ret?(h.m) \wedge inv?(h.n) \Rightarrow f.m < f.n \end{aligned}$$

Definition 9. We say OC is linearizable with respect to OA iff $OC \models_{OA} lin \wedge total$.

First, we show contextual trace refinement for a *canonical implementation* [16, 2, 17], which splits each sequential abstract operation call into three actions: an *invocation*, an *effect action* and a *response*.

Definition 10. For an abstract procedure $ph_t(\mathbf{val} \textit{ in}, \mathbf{res} \textit{ out}) = P_t$, the canonical implementation of the procedure is:

$$\begin{aligned} \neg dec.\widehat{pc}_t &\rightarrow \mathbf{var} \widehat{pc}_t; \widehat{pc}_t := 1; H \wedge \langle inv(t, p, in) \rangle \\ \sqcap \widehat{pc}_t = 1 &\rightarrow ph_t(in, out); \widehat{pc}_t := 2 \\ \sqcap \widehat{pc}_t = 2 &\rightarrow \mathbf{rav} \widehat{pc}_t; H \wedge \langle ret(t, p, out) \rangle \end{aligned}$$

Invocation and response actions modify the auxiliary history variable by recording the corresponding event, while the effect action has the same effect as the abstract operation call. Unlike the abstract object, the histories of a canonical implementation are potentially concurrent.

Theorem 2 (Canonical contextual trace refinement). Suppose OA and OB are objects, where OB is a canonical implementation of OA . Then $OA \sqsubseteq OB$.

Proof. We use Lemma 1 because OB may not satisfy minimal progress. Here, $rel.act.OB$ trivially satisfies (4) because by nature each procedure of a canonical object terminates. The proof of (3) requires further consideration because $rel.act.OB$ may deadlock. For example, OB may be a stack with a *pop* operation that blocks when the stack is empty. In such cases, because no data refinement is performed, the guard of the canonical object is false when the guard of the abstract object is false, allowing one to discharge (3). The remaining proof obligations are straightforward. \square

Next, we restate a completeness result by Schellhorn et al. [17], who have shown completeness of backward simulation for verifying linearizability. In particular, provided OC is a linearizable implementation of OA , they show that it is always possible to construct a backward simulation relation between the OC and the canonical implementation of OA .

Lemma 4 (Completeness of backward simulation [17]). Suppose OA, OB and OC are objects and $\mathcal{M}[OA]$ is continuous. If $OC \models_{OA} lin \wedge total$ and OB is a canonical implementation of OA , then there exists a total relation R such that both (5) and (6) hold between $\mathcal{M}[OB]$ and $\mathcal{M}[OC]$.

Finally, we prove our main result for linearizability, i.e., that linearizability and minimal progress together preserves contextual trace refinement.

Theorem 3. *Suppose object OC is linearizable with respect to OA , OC satisfies minimal progress, and $\mathcal{M}[OA]$ is continuous. If \mathcal{C} is a client such that $\mathcal{C}[OA]$ is continuous then $\mathcal{C}[OA] \sqsubseteq \mathcal{C}[OC]$.*

Proof. Construct a canonical implementation OB of OA . By transitivity of \sqsubseteq , the proof holds if both (a) $\mathcal{C}[OA] \sqsubseteq \mathcal{C}[OB]$ and (b) $\mathcal{C}[OB] \sqsubseteq \mathcal{C}[OC]$. Condition (a) holds by Theorem 2, and (b) holds by Theorem 1 (part 2), followed by Lemma 4. Application of Theorem 1 (part 2) is allowed because if $\mathcal{C}[OA]$ is continuous then $\mathcal{C}[OB]$ is continuous, whereas application of Lemma 4 is allowed because if R satisfies (5) and (6) for $\mathcal{M}[OB]$ and $\mathcal{M}[OC]$, then R also satisfies (5) and (6) for $\mathcal{C}[OB]$ and $\mathcal{C}[OC]$. \square

7.2 Sequential and quiescent consistency

We now consider contextual trace refinement for concurrent objects that satisfy sequential consistency and quiescent consistency, both of which are weaker than linearizability. Both conditions are total [7]. Additionally, sequential consistency disallows reordering of operation calls within a thread (see sc below), while quiescent consistency (see qc below) disallows reordering across a quiescent point (defined by qp below).

$$\begin{aligned} sc(h, f) &\hat{=} \forall m, n : \text{dom}(f) \bullet m < n \wedge \pi_1.(h.m) = \pi_1.(h.n) \wedge \\ &\quad \text{ret?}(h.m) \wedge \text{inv?}(h.n) \Rightarrow f.m < f.n \\ qp(m, h) &\hat{=} \forall n : \text{dom}(h) \bullet n \leq m \Rightarrow \neg \text{pi}(n, h[0..m]) \\ qc(h, f) &\hat{=} \forall m, k, n : \text{dom}(f) \bullet m < k < n \wedge qp(k, h) \Rightarrow f.m < f.n \end{aligned}$$

Definition 11. *An object OC is sequentially consistent with respect to OA iff $OC \models_{OA} sc \wedge total$, and OC is quiescent consistent with respect to OA iff $OC \models_{OA} qc \wedge total$.*

Our results for sequential consistency and quiescent consistency are negative — neither condition guarantees trace refinement of the underlying clients, regardless of whether the client program in question is *data independent*, i.e., the state spaces of the client threads outside the shared object are pairwise disjoint.

Theorem 4. *Suppose object OC is sequentially consistent with respect to object OA . Then it is not necessarily the case that $OA \hat{=} OC$ holds.*

Proof. Consider the program in Fig. 3, where the client threads are data independent — x is local to thread 1, while y and z are local to thread 2 — and s is assumed to be sequentially consistent. Suppose thread 1 is executed to completion, and then thread 2 is executed to completion. Because s is sequentially consistent, the first `pop` (at T3) may set x to 1, the second (at U2) may set y to 2. This gives the execution:

$\langle (x, y, z) \mapsto (0, 0, 0), (x, y, z) \mapsto (1, 0, 0), (x, y, z) \mapsto (1, 0, 1), (x, y, z) \mapsto (1, 2, 1) \rangle$
that cannot be generated when using the abstract stack AS from Fig. 1 for s . \square

Theorem 4 differs from the results of Filipović et al. [9], who show that for data independent clients, sequential consistency implies observational refinement. In essence, their result holds because observational refinement only considers the initial and final states of a client program — the intermediate states of a client’s execution are ignored. Thus, internal reorderings due to sequentially consistent objects have no effect when

```

Init x, y, z = 0;
Thread 1 ==      Thread 2 ==
T1: s.push(1);   U1: z := 1;
T2: s.push(2);   U2: s.pop(y);
T3: s.pop(x);

```

Fig. 3. Counter example for contextual trace refinement and sequential consistency

```

Init x, y, z = 0;
Thread 1 ==      Thread 2 ==
T1: s.push(1);   U1: s.pop(z)
T2: s.push(2);
T3: s.pop(x);
T4: s.pop(y);
T5: s.push(3);

```

Fig. 4. Counter example for contextual trace refinement and quiescent consistency

only observing pre/post states. One can develop hiding conditions so that observational refinement becomes a special case of contextual trace refinement, allowing one to obtain the result by Filipović et al [9]. Further development of this theory is left for future work. We now give our result for quiescent consistency.

Theorem 5. *Suppose object OC is quiescent consistent with respect to object OA . Then it is not necessarily the case that $OA \hat{\sqsubseteq} OC$ holds.*

Proof. Consider the program Fig. 4, where the client threads are data independent — x and y are local to thread 1, while z is local to thread 2 — and s is a quiescent consistent stack. The concrete program may generate the following observable trace:

$$\langle (x, y, z) \mapsto (0, 0, 0), (x, y, z) \mapsto (1, 0, 0), (x, y, z) \mapsto (1, 2, 0), (x, y, z) \mapsto (1, 2, 3) \rangle$$

Note that the *pop* operations at T3 and T4 have been reordered, which could happen if the execution of *pop* at U1 overlaps with T1, T2, T3 and T4. The trace above is not possible when the client uses the abstract stack AS from Fig. 1. \square

8 Conclusions

In this paper, we have developed a framework, based on action systems with procedures, for studying the link between the correctness conditions for concurrent objects and contextual trace refinement, which guarantees substitutability of objects within potentially non-terminating reactive clients. Thus, we bring together the previously disconnected worlds of correctness for concurrent objects and trace refinement within action systems. We have shown that linearizability and minimal progress together ensure contextual trace refinement, but sequential consistency and quiescent consistency are inadequate for guaranteeing contextual trace refinement regardless of whether clients communicate outside the concurrent object. The sequential consistency result contrasts earlier results for observational refinement, where sequential consistency is adequate when clients only communicate through shared objects [9].

We have derived the sufficient conditions for contextual trace refinement using the proof obligations for forwards and backward simulation. However, neither of these conditions have been shown to be necessary, leaving open the possibility of using weaker correctness conditions on the underlying concurrent objects. Studying this relationship remains part of future work — areas of interest include the study of how the correctness conditions for safety of concurrent objects under relaxed memory models [7] can be

combined with different scheduler implementations for progress (e.g., extending [15, 13]) to ensure contextual trace refinement.

Acknowledgements We thank John Derrick and Graeme Smith for helpful discussions. Brijesh Dongol is supported by EPSRC grant EP/N016661/1. “Verifiably correct high-performance concurrency libraries for multi-core computing systems”.

References

1. R. J. R. Back and J. von Wright. Trace refinement of action systems. In *CONCUR '94: Proceedings of the Concurrency Theory*, pages 367–384. Springer-Verlag, 1994.
2. R. Colvin, S. Doherty, and L. Groves. Verifying concurrent data structures by simulation. *Electr. Notes Theor. Comput. Sci.*, 137(2):93–110, 2005.
3. W. P. de Roever and K. Engelhardt. *Data Refinement: Model-oriented proof methods and their comparison*. Cambridge Tracts in Theor. Comp. Sci. Cambridge Univ. Press, 1996.
4. J. Derrick, G. Schellhorn, and H. Wehrheim. Mechanically verified proof obligations for linearizability. *ACM Trans. Program. Lang. Syst.*, 33(1):4, 2011.
5. S. Doherty. Modelling and verifying non-blocking algorithms that use dynamically allocated memory. Master’s thesis, Victoria University of Wellington, 2003.
6. B. Dongol. *Progress-based verification and derivation of concurrent programs*. PhD thesis, The University of Queensland, 2009.
7. B. Dongol, J. Derrick, G. Smith, and L. Groves. Defining correctness conditions for concurrent objects in multicore architectures. In J. T. Boyland, editor, *ECOOP*, volume 37 of *LIPICs*, pages 470–494. Dagstuhl, 2015.
8. B. Dongol and L. Groves. Towards linking correctness conditions for concurrent objects and contextual trace refinement. REFINE workshop, 2015. To appear.
9. I. Filipović, P. W. O’Hearn, N. Rinetzky, and H. Yang. Abstraction for concurrent objects. *Theor. Comput. Sci.*, 411(51-52):4379–4398, 2010.
10. A. Gotsman and H. Yang. Liveness-preserving atomicity abstraction. In L. Aceto, M. Henzinger, and J. Sgall, editors, *ICALP(2)*, volume 6756 of *LNCS*, pages 453–465, 2011.
11. J. He, C. A. R. Hoare, and J. W. Sanders. Data refinement refined. In B. Robinet and R. Wilhelm, editors, *ESOP*, volume 213 of *LNCS*, pages 187–196. Springer, 1986.
12. M. Herlihy and N. Shavit. *The Art of Multiprocessor Programming*. Morg. Kauf., 2008.
13. M. Herlihy and N. Shavit. On the nature of progress. In A. Fernández Anta, G. Lipari, and M. Roy, editors, *OPODIS*, volume 7109 of *LNCS*, pages 313–328. Springer, 2011.
14. M. P. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.*, 12(3):463–492, 1990.
15. H. Liang, J. Hoffmann, X. Feng, and Z. Shao. Characterizing progress properties of concurrent objects via contextual refinements. In P. R. D’Argenio and H. C. Melgratti, editors, *CONCUR*, volume 8052 of *LNCS*, pages 227–241. Springer, 2013.
16. N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.
17. G. Schellhorn, J. Derrick, and H. Wehrheim. A sound and complete proof technique for linearizability of concurrent data structures. *ACM TOCL*, 15(4):31:1–31:37, 2014.
18. K. Sere and M. A. Waldén. Data refinement of remote procedures. *Formal Asp. Comput.*, 12(4):278–297, 2000.
19. R. K. Treiber. Systems programming: Coping with parallelism. Technical Report RJ 5118, IBM Almaden Res. Ctr., 1986.