

Number 808



UNIVERSITY OF  
CAMBRIDGE

Computer Laboratory

## Resource-sensitive synchronisation inference by abduction

Matko Botinčan, Mike Dodds,  
Suresh Jagannathan

January 2012

15 JJ Thomson Avenue  
Cambridge CB3 0FD  
United Kingdom  
phone +44 1223 763500  
<http://www.cl.cam.ac.uk/>

© 2012 Matko Botinčan, Mike Dodds, Suresh Jagannathan

Technical reports published by the University of Cambridge  
Computer Laboratory are freely available via the Internet:

*<http://www.cl.cam.ac.uk/techreports/>*

ISSN 1476-2986

# Resource-Sensitive Synchronization Inference by Abduction

Matko Botinčan      Mike Dodds      Suresh Jagannathan

## Abstract

We present an analysis which takes as its input a sequential program, augmented with annotations indicating potential parallelization opportunities, and a sequential proof, written in separation logic, and produces a correctly-synchronized parallelized program and proof of that program. Unlike previous work, ours is not an independence analysis; we insert synchronization constructs to preserve relevant dependencies found in the sequential program that may otherwise be violated by a naïve translation. Separation logic allows us to parallelize fine-grained patterns of resource-usage, moving beyond straightforward points-to analysis.

Our analysis works by using the sequential proof to discover dependencies between different parts of the program. It leverages these discovered dependencies to guide the insertion of synchronization primitives into the parallelized program, and to ensure that the resulting parallelized program satisfies the same specification as the original sequential program, and exhibits the same sequential behaviour. Our analysis is built using frame inference and abduction, two techniques supported by an increasing number of separation logic tools.

## 1 Introduction

Automatically verifying the safety properties of shared-memory concurrent programs remains a challenging problem. To be useful in practice, proof tools must (a) explore a potentially large number of interleavings, and (b) construct precise flow- and path-sensitive abstractions of a shared heap.

Just as significantly, verification complexity is often at odds with the straightforward intentions of the programmer. Low-level concurrency control abstractions such as locks obfuscate higher-level notions such as atomicity and linearizability, likely to be exploited by the programmer when writing programs. While attempts to incorporate these notions directly into programs have met with some success—for example, in software transactional memory [20], there remains a substantial burden on the programmer to ensure programs can be partitioned into easily-understood atomic and thread-local sections. Automated proof engines used to verify program correctness must also assimilate this knowledge to ensure proof-search focusses on relevant behaviour (e.g., serializability), and eschews irrelevant details (e.g., thread-local state) [17]; the alternative, expressing such distinctions in programmer-supplied specifications, is too often onerous to be acceptable in practice. In any case, the inherent complexity of dealing with concurrency, both for the programmer and for the verifier, unfortunately remains.

In many cases concurrency is an optimization, rather than intrinsic to the behaviour of the program. That is, a concurrent program is often intended to achieve the same effect of a simpler sequential counterpart. Consequently, an attractive alternative to *constructing* a concurrent programming is to automatically *synthesize* one. In this approach, the programmer writes a sequential program; the program is automatically transformed into a concurrent one exhibiting the same behaviour [4, 3, 5, 16]<sup>1</sup>. In this context, the verification problem becomes significantly more tractable. Understanding and verifying the concurrent program reduces to first verifying the sequential program, and second, verifying the parallelising transformation.

We propose a program analysis and transformation that yields a parallelized program, guided by a safety proof of the original sequential program. To guide the transformation, we assume that the programmer annotates points in the sequential program where concurrency can be profitably exploited, *without* supplying any additional concurrency control or synchronization. The result of the transformation is a concurrent program with corresponding behaviour, and a safety proof for the concurrent program; in this sense, parallelized programs are *verified by construction*. While our analysis is driven by a safety proof, parallelization protects all behaviours, not just those specified in the proof. Thus we can apply our approach to complex algorithms without needing to verify total functional correctness.

Our transformation ensures that behaviour is preserved by requiring that the concurrent program respects sequential data-dependencies. In other words, the way threads access and modify shared resources never results in behaviour that would not be possible under sequential evaluation. To enforce these dependencies, the transformation injects *barriers*, signalling operations that regulate when threads are allowed to read or write shared state. These barriers can be viewed as resource transfer operations which acquire and relinquish access to shared resources such as shared-memory data structures and regions when necessary.

Our analysis is built on separation logic [31]. By tracking how resources are demanded and consumed within a separation logic proof, we can synthesize barriers to precisely control access to these resources. Our approach critically relies on frame inference and abduction [9], two techniques that generate fine-grained information on when resources are necessary and when they are redundant. This information enables optimizations that depend on deep structural properties of the resource—for example, we can split a linked list into dynamically-sized segments and transmit portions between threads piece-by-piece.

Our analysis thus enforces sequential order over visible behaviours, while allowing parallelization where it has no effect on behaviour. Insofar as our technique safely transforms a sequential program into a concurrent one, it can also be viewed as a kind of proof-directed compiler optimization. Notably, our transformation is *not* based on an independence analysis—our focus is not to identify when two potentially concurrent code regions do not interfere; instead, our analysis injects sufficient synchronization to ensure relevant sequential dependencies are preserved, even when regions share state and thus potentially interfere.

## Contributions.

---

<sup>1</sup>This approach is commonly known as *deterministic parallelism*, although our approach does not, in fact, require that the input sequential program is deterministic.

1. We present an automated technique to synthesize a parallel program given a partially-specified sequential program augmented with annotations indicating computations that are candidates for parallelization. The provided specifications are used to define relevant loop invariants.
2. We leverage abduction and frame inference to define a path- and context-sensitive program analysis capable of identifying per program-point resources that are both *redundant*—these are resources that would no longer be used by the thread executing this computation; and, *needed*—these are resources that would be required by any thread that executes this computation, but which is not known to have been released at this point.
3. We use information from the above analysis to inject *grant* and *allowed* barriers into the original program; their semantics enable resource transfer from redundant to needed resource points. The analysis also constructs a safety proof in separation logic which validates the correctness of the barrier injections.
4. We prove that our transformation is specification-preserving: the parallelized program is guaranteed to satisfy the same specification as the sequential program. Moreover, for terminating programs that do not dispose memory, we also show that the transformed program preserves the *behaviour* of the original. Complete proofs and other supporting material can be found in appendices.

**Paper structure.** An overview of the approach and motivating examples are given in §2. Issues that arise in dealing with loops and recursive datatypes (such as lists) are discussed in §3. The analysis and transformation are formalized in §4 and §5. Behaviour preservation results are discussed in §6. Observations about the interplay between analysis precision and the predicates available in defining specifications are given in §7. §8 discusses related work.

## 2 Overview

The objective of our analysis is to take a sequential program (with annotations indicating where concurrency may be exploited), and to produce a *parallelized* program that conforms to the same specification. To do this, we insert *synchronization barriers* enforcing those sequential dependencies that can affect the observable behaviour of the program.

Our analysis uses *separation logic* to capture and manipulate these dependencies. The programmer must provide a proof of safety for the sequential program. This proof is used to drive the parallelization process, allowing our analysis to calculate precisely which resources can be accessed in parallel, and which must be accessed sequentially, in order to preserve salient sequential dependencies.

### 2.1 Parallelization and Parallel-for

We assume that the portions of a program that can be parallelized will be written as *parallel-for* loops, `pfor`. The intended semantics of a parallel-for loop is that every iteration will run in parallel, but that the behaviour will be identical to running them in

```

global *x, *y;

void main() {
    local i, n;
    n = nondet();
    x = alloc();
    y = alloc();
    pfor(i=1;i++;i<n){
        f(i);
    }
}

void f(i){
    local v=*x;
    if (v>=i) {
        g(y, v);
    }
    else {
        g(x, 0);
    }
}

void g(*p, v){
    *p = v;
}

```

Figure 1: A simple parallel program that operates over shared locations `x` and `y`. The number of concurrent iterations performed is chosen nondeterministically to represent the act of receiving values from some unknown process or user interaction.

---

sequence; this semantics provides a simple but useful form of data parallelism in which concurrently executing computations may nonetheless have access to shared state.

For example, consider the program shown in Fig. 1. How can we parallelize this program without introducing any unwanted new behaviours, i.e., behaviours not possible under sequential execution? Naïvely, we might simply run every iteration of the `pfor` in parallel, without synchronization. That is:

```
f(1) || f(2) || ...
```

In some cases, this parallelization is good, and introduces no new observable behaviours (e.g., if `v` is greater than `n`). But, under others, the second call to `f()` may write to a memory location that is subsequently read by the first call to `f()`, violating intended sequential ordering. For example, consider the case when `v` is initially 1 and `n` is 2; sequential execution would produce a final result in which the locations pointed to by `x` and `y` resp. are 0 and 1, while executing the second iteration before the first would yield a result in which the location pointed to by `y` remains undefined.

To be sure that no new observable behaviour is introduced, we must ensure that:

- no iteration can read from a location that was already written to by a later iteration.
- no iteration can write to a location that was already read from or written to by a later iteration.

In order to enforce this behaviour, we must introduce a mechanism to force later iterations to wait for earlier ones.

Note, crucially, that later iterations *need not always* wait for earlier ones—synchronization is only needed when reads and writes to a particular memory location could result in out-of-order behaviour. We want to enforce *salient* dependencies, while allowing beneficial race-free concurrency.

## 2.2 Dependency-Enforcing Barriers

Our analysis inserts barriers into the parallelized program requiring logically later iterations to wait for earlier ones. We introduce `grant()`, a barrier that signals to subsequent

```

f1(i){
  local v=*x;
  if (v>=i) {
    grant(wx);
    ga1(y, v);
  }
  else {
    grant(wy);
    gb1(x, 0);
  }
}

ga1(*p, v){
  *p=v;
  grant(wy);
}

gb1(*p, v){
  *p=v;
  grant(wx);
}

f2(i){
  wait(wx);
  local v=*x;
  if (v>=i) {
    ga2(y, v);
  }
  else {
    gb2(x, 0);
    wait(wy);
  }
}

ga2(*p, v){
  wait(wy);
  *p=v;
}

gb2(*p, v){
  *p=v;
}

```

Figure 2: Parallelization of  $f()$ . Only synchronization barriers between the first and second iterations are shown.

---

iterations that a resource can safely be accessed, and `wait()`, a barrier that blocks until the associated resource becomes available from preceding iterations [16].

How can we use these barriers to enforce sequential dependencies in the example program? The exact pattern of parallelization depends on the granularity of our parallelization analysis. In the best case, there are no dependencies between iterations (e.g., each iteration operates on a different portion of a data structure). In this case, we need no barriers, and all iterations run independently. However, our example program shares the locations  $x$  and  $y$ , meaning barriers are required. In the worst case, each iteration of the `pfor` begins with a call to `wait()` and ends with a call to `grant()`—this would enforce sequential ordering on the iterations.

A better (although still relatively simple) parallelization is shown in Fig. 2. To simplify the presentation, we only show the synchronization barriers between the first and second iterations of `pfor`. We write  $f_1$  for the first iteration, and  $f_2$  for the second. (Our full analysis generates a single version of  $f()$  with sufficient barriers to run an arbitrary number of iterations in parallel. See §2.6 for a description.)

Two channels,  $wx$  and  $wy$ , are used to signal between  $f_1$  and  $f_2$ — $wx$  (resp.  $wy$ ) is used to signal that the later thread can read and write to the heap location referred to by  $x$  (resp.  $y$ ). The function  $g()$  has different resource requirements depending on its calling context. Consequently, we have specialized it to embed appropriate barriers depending on its context.

Our analysis inserts barriers so that an iteration of the parallel-for only calls `grant()` on a signal when the associated resource will no longer be accessed. Similarly, it ensures that an iteration calls `wait()` to acquire the resource before the resource is accessed.

How does our analysis generate these synchronization barriers? The example we have given here is simple, but in general our analysis must cope with complex dynamically-allocated resources such as linked lists. It must deal with the partial ownership (for example, read-access), and with manipulating portions of a dynamic structure (for example, just the head of a linked list). We therefore need a means to express complex (dynamic) patterns of resource management and transfer.

To achieve this, our analysis assumes a sequential *proof*, written in separation logic, rather than just an undecorated program. This proof need not capture full functional correctness—it is sufficient just to prove memory-safety. We exploit the dependencies expressed within this proof to determine the resources that are needed at a program point, and when they can be released. Our analysis inserts barriers to enforce the sequential dependencies represented in this proof. As our proof system is sound, these dependencies faithfully represent those in the original sequential program.

## 2.3 Resources, Separation and Dependency

At the heart of our analysis is automated reasoning using concurrent separation logic [31, 27]. Separation logic is a Hoare-style program logic for reasoning about mutable, allocatable resources. A separation logic proof of a program  $C$  establishes a specification

$$\{P\} C \{Q\}$$

This can be read as saying: “the program  $C$ , if run in a resource satisfying  $P$ , will not fault, and will give a resource satisfying  $Q$  if it terminates”. In general, a *resource* can be anything for which ownership can be partitioned (‘separated’) between different threads [11]. In practice, resources are most often heap-allocated structures such as lists, trees, locks, etc, where separation corresponds to disjointness between underlying heap addresses.

An essential feature of separation logic is that specifications are *tight*. This means that all of the resources accessed by a program  $C$  must be described in its precondition  $P$ , or acquired through explicit resource transfer. No other resources will be accessed or affected by  $C$  when executed from a resource satisfying  $P$ . The tight interpretation is essential for our analysis. Suppose we prove a specification for a program (or portion thereof); resources outside of the precondition cannot affect the program’s behaviour, and consequently can be safely transferred to other threads.

One result of this tight interpretation is the *frame rule*, which allows a small specification to be embedded into a larger context.

$$\frac{\{P\} C \{Q\}}{\{P * F\} C \{Q * F\}} \text{FRAME}$$

Here  $*$  is the *separating conjunction*. An assertion  $P * F$  is satisfied if  $P$  and  $F$  hold, and the portions of the state they denote do not overlap.

The concurrent counterpart of the frame rule is the *parallel rule*. This allows two threads that access non-overlapping resources to run in parallel, without affecting each



other’s behaviour.

$$\frac{\{P_1\} C_1 \{Q_1\} \quad \{P_2\} C_2 \{Q_2\}}{\{P_1 * P_2\} C_1 || C_2 \{Q_1 * Q_2\}} \text{ PARALLEL}$$

The parallel rule enforces the absence of data-races between  $C_1$  and  $C_2$ . The two threads can consequently only affect each other’s behaviour by communicating through race-free synchronization mechanisms, such as locks or barriers.

**Automatic inference.** Automated reasoning in separation logic revolves around two kinds of inference questions. The first, *frame inference*, calculates the portion of a formula  $S$  which is ‘left over’, once another formula  $P$  has been satisfied. We call this remainder  $F_?$  the *frame*, and write the inference question as follows:

$$S \vdash P * [F_?]$$

(Throughout the paper, we use square brackets to indicate the portion of the entailment that is to be computed.)

The second kind of inference question, *abduction*, calculates the ‘missing’ formula that must be combined with a formula  $S$  in order to satisfy some other formula  $P$ . We call this formula  $A_?$  the *antiframe*, and write the inference question as follows:

$$S * [A_?] \vdash P$$

Frame inference and abduction form the basis of symbolic execution in separation logic [2, 9]. Intuitively, frame inference lets us reason forwards, while abduction lets us reason backwards. Frame inference can work out which portions of the state will (and will not) be affected by the command, while abduction can work out what extra resources are necessary to execute the command safely.

Suppose we have a symbolic state represented by an assertion  $S$ , and a command  $c$  with specification  $\{P\} c \{Q\}$ . If we calculate the frame in  $S \vdash P * [F_?]$ , the symbolic state after executing  $c$  must be  $Q * F_?$ . The tight interpretation of triples is necessary for this kind of reasoning. Because a specification must describe all the resources affected by a thread, any resources in the frame must be unaffected. Conversely, if we calculate the antiframe  $S * [A_?] \vdash P$ , it must be the case that before executing  $c$ , we must first acquire the additional resource  $A_?$  (as well as  $S$ ).

**Redundant and needed resources.** The tight interpretation means that a proof in separation logic expresses all resources modified by each individual command. Because separation logic ensures race-free behaviour, a proof also expresses all the resources that can affect the observable behaviour of each individual command in the program. Our parallelization analysis uses this to calculate resources that are redundant and those that are needed during the execution of the program.

We use frame inference to determine *redundant resources*—resources that will not be accessed in a particular portion of the program, and which can thus be safely transferred to other threads. Conversely, we use abduction to determine *needed resources*—resources that must be held for a particular portion of the program to complete, and which must thus be acquired before the program can proceed safely.

In our analysis, we generally calculate the redundant resource from the current program point to the start of the subsequent iteration. This resource can be transferred to the subsequent iteration using a `grant()` barrier. We generally calculate the needed resource from the end of the previous iteration to the current program point. This resource must be acquired from the previous iteration using a `wait()` barrier before execution can proceed further.

Note that these two resources need not be disjoint. A resource may be used early in an iteration, in which case it will be both needed from the previous iteration to the current point, and redundant from the current point up to the next iteration. Note also that redundant and needed resources need not cover the entire state—some resource described in the proof may never be accessed or modified by the program.

## 2.4 Algorithm Overview

The user supplies a sequential program which makes use of the `pfor` parallel-for construct, as well as a sequential proof written in separation logic establishing the memory-safety of the program. The high-level structure of our algorithm is then as follows:

1. The *resource usage analysis* phase uses abduction and frame inference to discover redundant and needed resources for different portions of the program.
2. The *parallelising transformation* phase consists of two parts:
  - (a) The *resource matching* phase matches redundant resources in one iteration of the parallel-for with needed resources in the subsequent iteration.
  - (b) The *barrier insertion* phase converts the sequential program into a concurrent program and inserts `grant()` and `wait()` barriers consistent with the discovered resource-transfer.

The result is a parallelized program, and its separation logic proof.

## 2.5 Resource Usage Analysis

The resource usage analysis takes as its input the program, and a separation logic proof for the program for some specification. Consider once again the functions `f()` and `g()`, which we introduced at the start of this section. Let us assume the programmer proves the following specification:

$$\{\mathbf{x} \mapsto x * \mathbf{y} \mapsto y\} \quad \mathbf{f}(i) \quad \{(x \geq i \wedge \mathbf{x} \mapsto x * \mathbf{y} \mapsto x) \vee (\mathbf{x} \mapsto 0 * \mathbf{y} \mapsto y)\}$$

Fig. 3 shows an outline proof of this specification.

For each program point, the resource usage analysis computes a pair of assertions: the *needed* resource and the *redundant* resource (see Fig. 4). The needed resource is the portion of the sequential precondition required to reach the program point without faulting; the redundant resource is the portion of the resource held in the sequential proof that is unnecessary for reaching the end of the function. In other words, needed assertions at program point  $p$  denote resources that are needed from the beginning of the function

<pre> {x ↦ x * y ↦ y} void f(i) {   {x ↦ x * y ↦ y}   local v = *x;   {v = x ∧ x ↦ x * y ↦ y}   if (v &gt;= i) {     {v = x ∧ v ≥ i ∧ x ↦ x * y ↦ y}     g(y, v);   }   else {     {v = x ∧ v &lt; i ∧ x ↦ x * y ↦ y}     g(x, 0);   } } {(x ≥ i ∧ x ↦ x * y ↦ x) ∨ (x ↦ 0 * y ↦ y)} </pre>	<pre> {p ↦ - ∧ v = v} void g(*p, v) {   {p ↦ - ∧ v = v}   *p = v; } {p ↦ v} </pre>
---	--

Figure 3: Proofs of functions  $f()$  and  $g()$ .

up to (but not including)  $p$ ; redundant assertions denote resources that are not required from  $p$  to the end of the function.

For example, in  $f$  the needed resource at the start of the **else**-branch of the conditional is:

$$x < i \wedge x \mapsto x$$

This asserts that the thread reaching this point still requires access to  $x$ , provided  $v < i$ . According to the sequential proof in Fig. 3, the program *actually* holds the resource:

$$v = x \wedge v < i \wedge x \mapsto x * y \mapsto y$$

To calculate the redundant resource, we ask the following frame-inference question, computing the frame  $F_?$ .

$$v = x \wedge v < i \wedge x \mapsto x * y \mapsto y \quad \vdash \quad x < i \wedge x \mapsto x * [F_?]$$

The resulting redundant resource is as follows:

$$F_?: \quad v = x \wedge x < i \wedge y \mapsto y$$

This asserts that a thread executing this function will no longer require access to  $y$  at this point, whenever  $v < i$ .

## 2.6 Generalizing to $n$ Threads

Our analysis can handle an unbounded number of loop iterations running in parallel. Fig. 5 shows the final parallelization of our example, generalized to deal with an arbitrary number of iterations.

A **pfor** is translated to a sequential **for** loop, in which each iteration forks a new copy of the parallelized loop body. Resources are transferred between the threads in order

```

void f(i) {
  n: {emp}
  r: {x < i ∧ y ↦ y}
  local v = *x;
  n: {x ↦ x}
  r: {x ≥ i ∧ x ↦ x} ∨ (x < i ∧ y ↦ y)
  if (v ≥ i) {
    n: {x ≥ i ∧ x ↦ x}
    r: {x ≥ i ∧ x ↦ x}
    g(y, v); → void g(*p, v) {
      n: {x ≥ i ∧ x ↦ x}
      r: {x ≥ i ∧ x ↦ x}
      *p = v;
      n: {x ≥ i ∧ x ↦ x * y ↦ y}
      r: {x ↦ x * y ↦ y}
    }
  }
  else {
    n: {x < i ∧ x ↦ x}
    r: {v = x ∧ x < i ∧ y ↦ y}
    g(x, 0); → void g(*p, v) {
      n: {x < i ∧ x ↦ x}
      r: {v = x ∧ x < i ∧ y ↦ y}
      *p = v;
      n: {x < i ∧ x ↦ x}
      r: {x ↦ x * y ↦ y}
    }
  }
  n: {(x ≥ i ∧ x ↦ x * y ↦ y) ∨ (x < i ∧ x ↦ x)}
  r: {x ↦ x * y ↦ y}
}

```

Figure 4: Redundant and needed resources in the function  $f()$ . Function  $g()$  is inlined, as its resource usage depends on the calling context.

of thread-creation. That is, the  $n$ th iteration of the `pfor` acquires resources from the logically-preceding  $(n - 1)$ th iteration, and releases resources to the logically  $(n + 1)$ th iteration. This ordering is implemented through shared channels—a thread shares with its predecessor a set of channels for receiving resources, and with its successor a set of channels for sending resources.

The number of threads—and so, the required number of channels—is potentially decided at run-time. Consequently, channels are dynamically allocated in the main `for`-loop using the `newchan()` operation [16]. Each iteration creates a set of new channels, and passes the prior and new set to the forked thread. The parallelized version of `f` (Fig. 5) now takes four channel arguments, a pair of each for `x` and `y`. The prior channels are used for resource transfer with the logically-preceding thread (here `wx`, `wy`), and the new channels are used to communicate resource transfer with the logically-following thread (here, `wx'`, `wy'`).

Our analysis generates two versions of function `g`, specialized to the two contexts in which it is invoked. Function `ga` executes when `v >= i`. Here, the executing thread needs write access to `y`, which it acquires by calling `wait(wyp)`. Function `gb` needs write access to `x`, but it runs in a context where the resource `x` has already been acquired. Consequently, it need not call `wait()`. Both versions release a resource to the following thread using `grant()` before returning.

### 3 Loops and Recursive Data Structures

Up to this point, we have dealt with channels transferring single heap locations injected into straight-line code. Our analysis can in fact handle more complicated heap-allocated data-structures such as linked lists, and control-flow constructs such as loops. To illustrate, consider the example program `sum_head(n)` shown (and verified) in Fig. 6. `sum_head(n)` sums the values stored in the first  $n$  elements of the list, then zeros the rest of the list.

An important feature of our approach is that the input safety proof need not specify all the relevant sequential properties; all sequential dependencies are enforced. Thus we can see `sum_head` as representative of the class of algorithms that traverse the head of a list, then mutate the tail. With minor modifications, the same proof pattern would cover insertion of a node into a list, or sorting the tail of the list.

Consider two calls to this function within a `parallel-for`:

```
void main(){
  pfor(i=1; i++; i<3){
    n=nondet(); sum_head(n);
  }
}
```

As above, for the sake of exposition we specialize the first and second iterations of the `pfor` loop, calling them `sum_head1` and `sum_head2`. Our analysis generalizes to  $n$  iterations exactly as described in §2.6. The barriers injected into `sum_head` must enforce the following properties:

```

void main() {
    local i, n = nondet();
    x = alloc(); y = alloc();
    chan wx', wy';
    chan wx = newchan(); chan wy = newchan();
    grant(wx); grant(wy);
    for (i = 0; i++; i < n) {
        wx' = newchan(); wy' = newchan();
        fork(f(i, wx, wy, wx', wy'));
        wx = wx'; wy = wy';
    }
    wait(wx); wait(wy);
}

f(i, wxp, wyp, wx, wy){
    wait(wxp);
    local v = *x;
    if (v >= i) {
        grant(wx);
        ga(y, v, wy);
    }
    else {
        wait(wyp);
        grant(wy);
        gb(x, 0, wx);
    }
} }

ga(*p, v, wyp, wy){
    wait(wyp);
    *p = v;
    grant(wy);
}

gb(*p, v, wx){
    *p = v;
    grant(wx);
}

```

Figure 5: Parallelization of our running example generalized to deal with  $n$  threads. Channels are used to signal availability of resources from one thread to another.

```

lnode *hd;
{hd ↦ h * lseg(h, nil)}
sum_head(n) {
  local i, sum, x;
  int i = 1;
  int sum = 0;
  x = *hd;
  {hd ↦ x * lseg(x, nil)}
  // entailment on lseg predicate
  {hd ↦ h * lseg(h, x) * lseg(x, nil)}
  while(x!=nil && i!=n) {
    {hd ↦ h * lseg(h, x) * ∃v, y. x.val ↦ v * x.nxt ↦ y * lseg(y, nil)}
    sum += x.val;
    i++;
    x = x.nxt;
  }
  {hd ↦ h * lseg(h, x) * lseg(x, nil)}
  while(x!=nil) {
    {hd ↦ h * lseg(h, x) * ∃v, y. x.val ↦ v * x.nxt ↦ y * lseg(y, nil)}
    x.val = 0;
    x = x.nxt;
  }
}
{hd ↦ h * lseg(h, nil)}

```

Figure 6: Separation logic proof of `sum_head`, a list-manipulating program whose automated parallelization requires reasoning over complex assertions and predicates.

```

sum_head1(n){
  i = 1;
  sum = 0;
  x = *hd;
  while(x!=nil && i!=n){
    sum+=x.val;
    i++;
    x = x.nxt;
  }
  *xpr = x;
  grant(i1);
  while(x!=nil){
    x.val = 0;
    x = x.nxt;
  }
  grant(i2);
}

sum_head2(n){
  i = 1;
  sum = 0;
  wait(i1);
  x = *hd;
  while(x!=nil && i!=n){
    if(x==*xpr) wait(i2);
    sum +=x.val;
    i++;
    x = x.nxt;
  }
  while(x!=nil){
    if(x==*xpr) wait(i2);
    x.val = 0;
    x = x.nxt;
  }
}

```

Figure 7: Parallelization of `sum_head` for two threads.

- `sum_head2` must not write to a list node until `sum_head1` has finished both writing to and reading from it. Consequently, if  $n_2 < n_1$ , `sum_head2` must wait for `sum_head1` to finish summing the values stored in the first  $n_1$  nodes before writing to them.
- `sum_head2` must not read from a list node until `sum_head1` has finished writing to them. Consequently, `sum_head2` must wait for `sum_head1` to finish zeroing the value stored in a node before reading from the node.

This example is substantially more subtle than our earlier one, requiring more than a simple points-to analysis, because the list is not divided into statically-apparent reachable segments. In the worst case, a `wait()` at the start of `sum_head2()` and a `grant()` at the end of `sum_head1()` enforces sequential order. However, by reasoning about the structure of the manipulated list using the safety proof given in Fig. 6, our approach can do considerably better.

The parallelization synthesized by our analysis is shown in Fig. 7 (the general  $k$ -thread version is shown in Fig. 8). This parallelization divides the list into two segments, consisting of the portions read and written to by `sum_head1()`. A shared heap-location `xpr` stores the starting address of the portion written by `sum_head1()`. The thread `sum_head2()` uses `xpr` to control when to access the second segment of the list<sup>2</sup>. We discuss how the analysis materializes `xpr` below.

Handling dynamic structures means dealing with allocation and disposal. Fortunately, separation logic handles both straightforwardly. Updates to the data-structure and object allocation are by assumption reflected in the invariants of the original sequential proof. Thus updates and allocations are also reflected in the invariants which our analysis constructs to represent the contents of channels. However, introducing allocation and

---

<sup>2</sup>For simplicity, here we write `sum_head2` with `wait` controlled by conditionals. The actual transformation performs syntactic loop-splitting to avoid the need to modify the loop invariant. Details are given in §5.2.



```

void main() {
  local i, k = nondet();
  x = alloc(); y = alloc();
  chan i1, i2;
  chan i1p = newchan();
  chan i2p = newchan();
  grant(i1p); grant(i2p);
  for (i = 0; i++; i < k) {
    xft = alloc();
    i1 = newchan(); i2 = newchan();
    fork(sum_head(n, i1p, i2p,
                  i1, i2, xpr, xft));
    xpr = xft;
    i1p = i1; i2p = i2;
  }
  wait(i1p); wait(i2p);
}

sum_head(n, i1p, i2p, i1, i2, xpr, xft){
  i = 1;
  sum = 0;
  x = *hd;
  wait(i1p);
  while(x!=nil && i!=n){
    if(x==*xpr) wait(i2p);
    sum+=x.val;
    i++;
    x = x.nxt;
  }
  *xft = x;
  grant(i1);
  while(x!=nil){
    if(x==*xpr) wait(i2p);
    x.val = 0;
    x = x.nxt;
  }
  grant(i2);
}

```

Figure 8: Generalised parallelisation of `sum_head` for  $k$  threads.

disposal affects the behaviour-preservation result discussed in §6; this result ensures the program behaviour is unaffected by the translation (i.e. the translation enforces deterministic parallelism).

**Reasoning about list segments.** We assume that our separation logic domain includes the predicate  $\text{lseg}(x, t)$ , which asserts that a segment of a linked list exists with head  $x$  and tail-pointer  $t$ . We define  $\text{lseg}$  as the least separation logic predicate satisfying the following recursive equation<sup>3</sup>:

$$\text{lseg}(x, t) \triangleq x = t \wedge \text{emp} \vee \exists v, y. x.\text{val} \mapsto v * x.\text{nxt} \mapsto y * \text{lseg}(y, t)$$

We assume that the programmer proves the following specification for the sequential version of `sum_head`:

$$\{\text{hd} \mapsto h * \text{lseg}(h, \text{nil})\} \quad \text{sum\_head}(n) \quad \{\text{hd} \mapsto h * \text{lseg}(h, \text{nil})\}$$

This specification is trivial: all it says is that executing `sum_head` beginning with a list segment, results in a list segment. Fig. 6 shows a proof of this specification.

We run our resource-usage analysis over the program to determine redundant and needed resources. Consider the following loop from the end of `sum_head1`:

---

<sup>3</sup>Our analysis depends strongly on the choice of these basic predicates. See §7 for a discussion of alternatives to  $\text{lseg}$ .

```

...
while(x!=nil) {
  x.val = 0; x = x.nxt;
}

```

Our analysis reveals that only the resource defined by  $\text{lseg}(\mathbf{x}, \text{nil})$  is needed from the start of this loop to the end of the iteration. Comparing this resource to the corresponding invariant in the sequential proof reveals that the resource  $\exists h. \text{hd} \mapsto h * \text{lseg}(h, \mathbf{x})$  is redundant at this point. This assertion represents the segment of the list that has already been traversed, from the head of the list to  $\mathbf{x}$ .

**Materialization and barrier injection.** Notice that the assertions generated by our analysis are partly expressed in terms of the local variable  $\mathbf{x}$ , which may change during execution. In order to safely transfer these assertions to subsequent iterations of the `pfor`, we need them to be *globally accessible* and *invariant*. To satisfy this goal, we could simply existentially quantify the offending variable,  $\mathbf{x}$ , giving a redundant invariant

$$\exists h, y. \text{hd} \mapsto h * \text{lseg}(h, y)$$

However, such a weakening loses important information, in particular the relationship between the necessary resource, the list segment from  $\mathbf{x}$  and the tail of the list. To retain such dependency relationships, our analysis *materializes* the current value of  $\mathbf{x}$  into a location `xpr` shared between `sum_head1` and `sum_head2`. An assignment is injected into `sum_head1` at the start of the second loop:

```

...
*xpr = x;
while(x!=nil) {
...

```

After the assignment to `xpr`, the redundant state can now be described as follows:

$$\exists h, y. \text{hd} \mapsto h * \text{lseg}(h, y) * \text{xpr} \xrightarrow{1/2} y$$

Here we use fractional permissions in the style of [6] to allow a location to be shared between threads. The assertion  $\text{xpr} \xrightarrow{1/2} y$  represents fractional, read-only permission on the shared location `xpr`. This helps in binding together the head of the list and the remainder of the list when they are recombined.

When traversing the list, `sum_head2` compares its current position with `xpr`. If it reaches the pointer stored in `xpr`, it must wait to receive the second, remainder segment of the list from `sum_head1`.

## 4 Technical Background

### 4.1 Programming Language and Representation

We assume the following heap-manipulating language<sup>4</sup>:

$$\begin{aligned}
e &::= \mathbf{x} \mid \mathbf{nil} \mid t(\bar{e}) \mid \dots && (\text{expressions}) \\
b &::= \mathbf{true} \mid \mathbf{false} \mid e = e \mid e \neq e \mid \dots && (\text{booleans}) \\
a &::= \dots && (\text{atomic commands}) \\
C &::= C; C \mid \ell: \mathbf{skip} \mid \ell: a \mid \ell: \mathbf{x} := f(\bar{e}) \mid \ell: \mathbf{return} e \\
&\quad \mid \ell: \mathbf{if}(b) \{ C \} \mathbf{else} \{ C \} \mid \ell: \mathbf{while}(b) \{ C \} \\
\mathbb{P} &::= \mathbf{global} \bar{r}; (f(\bar{x}) \{ \mathbf{local} \bar{y}; C \})^+
\end{aligned}$$

To simplify the exposition slightly, we assume a fixed input program  $\mathbb{P}$  of the following form:

$$\mathbf{main}_{\mathbb{P}}() \{ \mathbf{pfor}(C_1; C_2; b) \{ \mathbf{work}(); \} \}$$

Our parallelization approach generates a new program  $\mathbb{P}_{par}$  that allows safe execution of (a transformed version of)  $\mathbf{work}$  in separate threads. We denote by  $\mathbf{Func} \ni \{ \mathbf{main}, \mathbf{work} \}$  the set of functions declared in a program  $\mathbb{P}$ .

**Labelling of commands.** Every command  $C \in \mathbf{Cmd}$  in the program is indexed by a unique label  $\ell \in \mathbf{Label}$ ; function identifiers are also treated as labels. We identify a particular command by its label, and when needed denote by  $\mathbf{cmd}(\ell)$  the command at the label  $\ell$ . The set of all labels occurring in  $C$  is given by  $\mathbf{Lbs}(C)$ . The left-hand side of Fig. 9 shows a labelling of the functions  $\mathbf{f}$  and  $\mathbf{g}$ .

Commands within a block (function,  $\mathbf{if-else}$  or  $\mathbf{while}$ ) form a sequence identified by the sequence of the corresponding labels. Functions  $\mathbf{pred}, \mathbf{succ}: \mathbf{Label} \rightarrow \mathbf{Label}$  return the label of the previous (the next, resp.) command in the sequence. For  $\ell$  corresponding to a function name, a function call or a  $\mathbf{while}$  loop, the predecessor of the first command and successor of the last command in the block are denoted by  $\ell_s$  and  $\ell_e$ , respectively. For  $\ell$  corresponding to  $\mathbf{if-else}$ ,  $\ell_s^t$  ( $\ell_e^t$ ) and  $\ell_s^f$  ( $\ell_e^f$ ) are labels of the predecessor of the first (the successor of the last) commands in the  $\mathbf{if}$  and  $\mathbf{else}$  branches, respectively.

Given labels  $\ell$  and  $\ell'$  within the same block, we say that  $\ell < \ell'$  if  $\exists n \geq 1$  such that  $\ell' = \mathbf{succ}^n(\ell)$ . Let  $\ell_{\uparrow}$  and  $\ell_{\downarrow}$  denote the smallest (resp. largest) label in the block containing  $\ell$ . Let  $[\ell, \ell')$  denote the sequence of labels between  $\ell$  inclusively and  $\ell'$  exclusively, and let  $\mathbb{P}_{[\ell, \ell')}$  stand for the corresponding program fragment. Analogously, we define  $\langle \ell, \ell' \rangle$  and  $\mathbb{P}_{\langle \ell, \ell' \rangle}$ .

**Program representation.** A program is represented via a control-flow graph with  $\mathbf{Label}$  as the set of nodes and as the set of edges all  $(\ell, \ell')$  such that  $\ell$  corresponds to a label of a function call and  $\ell' \in \langle \ell_s, \ell_e \rangle$ , or  $\ell$  corresponds to  $\mathbf{if-else}$  and  $\ell' \in \langle \ell_s^t, \ell_e^t \rangle \cup \langle \ell_s^f, \ell_e^f \rangle$ . A *path delimiter* is a finite sequence of nodes in the control-flow graph. Intuitively, a path delimiter represents the sequence of function calls and conditionals that need to be traversed in order to reach a particular assertion. Due to the structure of our analysis, we are only interested in the set of path delimiters in the function  $\mathbf{work}$ , denoted by  $\mathbf{AP}$ .

<sup>4</sup>To avoid introducing an extra construct, we define  $\mathbf{for}(C_1; C_2; b)\{C_3\}$  as  $C_1; \mathbf{while}(b)\{C_2; C_3\}$ .

<pre> void f(int i){   ℓ<sub>1</sub>: int v = *x;   ℓ<sub>2</sub>: if (v &gt;= i){   ℓ<sub>3</sub>: g(y, v);       }       else{ ℓ<sub>4</sub>: g(x, 0);       } }  void g(int* p,       int v){ ℓ<sub>5</sub>: *p = v; } </pre>	<pre> f<sub>s</sub>: x ↦ x * y ↦ y ℓ<sub>1</sub>: x ↦ x * y ↦ y ℓ<sub>2</sub>: v = x ∧ x ↦ x * y ↦ y ℓ<sub>3</sub>: v = x ∧ v ≥ i ∧ x ↦ x * y ↦ y ℓ<sub>4</sub>: v = x ∧ v &lt; i ∧ x ↦ x * y ↦ y f<sub>e</sub>: (x ≥ i ∧ x ↦ x * y ↦ x) ∨       (x &lt; i ∧ x ↦ 0 * y ↦ y)  g<sub>s</sub>: p ↦ _ ∧ v = v ℓ<sub>5</sub>: p ↦ _ ∧ v = v g<sub>e</sub>: p ↦ v </pre>
--	--

Figure 9: Left: labels for commands in `f` and `g`. Right: associated assertions in the sequential proof.

We often want to manipulate and compare path delimiters. For  $\gamma = \ell_1 \dots \ell_n \in \mathbf{AP}$ , we write  $\gamma[i]$  to denote  $\ell_i$ ,  $\gamma[i..j]$  to denote  $\ell_i \dots \ell_j$  and  $|\gamma|$  to denote the length  $n$ . For path delimiters  $\gamma$  and  $\gamma'$ ,  $\gamma \wedge \gamma'$  denotes their longest common prefix, i.e., for  $k = |\gamma \wedge \gamma'|$  we have  $\forall j \leq k, \gamma \wedge \gamma' = \gamma[j] = \gamma'[j]$  and if  $|\gamma|, |\gamma'| > k$  then  $\gamma[k+1] \neq \gamma'[k+1]$ . We define a partial order  $\prec$  on  $\mathbf{AP}$  as follows. We say that  $\gamma \prec \gamma'$  iff  $\gamma \wedge \gamma' = \gamma$  and  $|\gamma| < |\gamma'|$ , or for  $k = |\gamma \wedge \gamma'|$  we have  $|\gamma|, |\gamma'| > k$  and  $\gamma[k+1] \prec \gamma'[k+1]$ . We say that  $\gamma \preceq \gamma'$  iff  $\gamma \prec \gamma'$  or  $\gamma = \gamma'$ .

**Lemma 1.**  $(\mathbf{AP}, \preceq)$  is a lattice with the least element  $\mathbf{work}_s$  and the greatest element  $\mathbf{work}_e$ .

For  $\Gamma \subseteq_{fin} \mathbf{AP}$  we define  $\max(\Gamma)$  as  $\max(\Gamma) := \{\gamma \in \Gamma \mid \neg \exists \gamma' \in \Gamma. \gamma \prec \gamma'\}$ . We define  $\min(\Gamma)$  analogously.

## 4.2 Assertion Language and Theorem Prover

Assertions in our approach are expressed using a class of separation logic formulae called *symbolic heaps*. A symbolic heap  $\Delta$  is a formula of the form  $\exists \bar{x}. \Pi \wedge \Sigma$  where  $\Pi$  (the pure part) and  $\Sigma$  (the spatial part) are defined by:

$$\begin{aligned} \Pi &::= \mathbf{true} \mid e = e \mid e \neq e \mid p(\bar{e}) \mid \Pi \wedge \Pi \\ \Sigma &::= \mathbf{emp} \mid x \mapsto e \mid \mathbf{lseg}(e, e) \mid s(\bar{e}) \mid \Sigma * \Sigma \end{aligned}$$

Here  $\bar{x}$  are logical variables,  $e$  ranges over expressions,  $p(\bar{e})$  is a family of pure (first-order) predicates (such as e.g., arithmetic inequalities, etc), and  $s(\bar{e})$  a family of other spatial predicates (such as e.g., doubly-linked lists, trees, etc) besides points-to and linked-list predicates previously discussed. We refer to the pure and the spatial part of  $\Delta$  as  $\Delta^\Pi$  and  $\Delta^\Sigma$ . We denote set of all quantifier-free first-order formulae built in the same way as  $\Pi$  but also allowing the  $\vee$  connective by  $\Pi_\vee$ .

During our analysis, we often need to substitute variables, for example when recasting an assertion into a different calling context. If  $\varrho = \bar{x} \mapsto \bar{e}$  is a mapping from variables in

$\bar{x}$  to  $\bar{e}$  then by  $\Delta[\rho]$  we denote the formula obtained by simultaneously substituting every occurrence of  $x_i$  in  $\Delta$  with the corresponding  $e_i$ . We denote by  $\delta^{-1}$  the inverse variable mapping, if  $\delta$  is injective<sup>5</sup>.

The set of all symbolic heaps is denoted by **SH**. We represent a disjunction of symbolic heaps as a set and interchangeably use the  $\cup$  and  $\vee$  operators. The set of all disjunctive symbolic heaps is  $\mathcal{P}(\text{SH})$ . We overload the  $\wedge$  and  $*$  operators in a natural way: for  $\Delta_i = \Pi_i \wedge \Sigma_i$ ,  $i = 1, 2$ , we define  $\Delta_1 * \Delta_2 = (\Pi_1 \wedge \Pi_2) \wedge (\Sigma_1 * \Sigma_2)$ ,  $\Pi \wedge \Delta_i = (\Pi \wedge \Pi_i) \wedge \Sigma_i$  and  $\Sigma * \Delta_i = \Pi_i \wedge (\Sigma * \Sigma_i)$ . Operators  $\wedge$  and  $*$  distribute over  $\vee$ , thus we allow these operations on disjunctive heaps just as if they were on symbolic heaps and furthermore use the same notation  $\Delta$  to refer to both symbolic and disjunctive symbolic heaps.

We assume a sufficiently powerful (automated) prover for separation logic that can deal with three types of inference queries:

- $\Delta_1 \vdash \Delta_2 * [\Delta_F]$  (frame inference): given  $\Delta_1$  and  $\Delta_2$  find the frame  $\Delta_F$  such that  $\Delta_1 \vdash \Delta_2 * \Delta_F$  holds;
- $\Delta_1 * [\Delta_A] \vdash \Delta_2$  (abduction): given  $\Delta_1$  and  $\Delta_2$  find the “missing” assumption  $\Delta_A$  such that  $\Delta_1 * \Delta_A \vdash \Delta_2$  holds;
- $\Delta_1 * [\Delta_A] \vdash \Delta_2 * [\Delta_F]$  (bi-abduction): given  $\Delta_1$  and  $\Delta_2$  find  $\Delta_A$  and  $\Delta_F$  such that  $\Delta_1 * \Delta_A \vdash \Delta_2 * \Delta_F$  holds.

As before, square brackets denote the portion of the entailment that should be computed. We sometimes write  $[-]$  for a computed assertion that is existentially quantified and will not be reused.

None of these queries has a unique answer in general. However, for our analysis any answer is acceptable (though some will give rise to a better parallelization than the others). Existing separation logic tools generally provide only one answer.

### 4.3 Sequential Proof

We assume a separation logic proof of the program  $\mathbb{P}$ , represented as a map  $\mathfrak{P}: \text{Label} \rightarrow \mathcal{P}(\text{SH})$ . Intuitively, assertions in the proof have the property that for any label  $\ell$  executing the program from a state satisfying  $\mathfrak{P}(\ell)$  up to some subsequent label  $\ell'$  will result in a state satisfying  $\mathfrak{P}(\ell')$ , and will not fault.

More formally, we assume functions  $\text{Pre}, \text{Post}: \text{Label} \rightarrow \mathcal{P}(\text{SH})$  associating labels of atomic commands and function calls with their pre- and post-condition, respectively, and a function  $\text{Inv}: \text{Label} \rightarrow \mathcal{P}(\text{SH})$  associating **while** labels with loop invariants. We also assume  $\Omega$ , a mapping from labels to variable substitutions such that  $\Omega(l) = \bar{x} \mapsto \bar{x}'$  maps formal variables  $\bar{x}$  in the specification assertion to actual variables  $\bar{x}'$  in the proof assertion. Finally, we assume a function  $\mathfrak{F}: \text{Label} \rightarrow \mathcal{P}(\text{SH})$  giving the framed portion of the state. We write  $\Delta[\Omega(l)]$  to represent the heap constructed by applying the substitutions defined by  $\Omega(l)$  to the assertion  $\Delta$ .

---

<sup>5</sup>In our framework, substitutions are always guaranteed to be injective because the variables being substituted correspond to heap locations and channel resources whose denotations are guaranteed to be distinct; if the substitution involves values, then they must be implicitly existentially quantified, and can therefore be assumed to be distinct.

Then at each label  $\ell$  we have  $\mathfrak{P}(\ell) \vdash \text{Pre}(\ell)[\Omega(\ell)] * \mathfrak{F}(\ell)$  and  $\text{Post}(\ell)[\Omega(\ell)] * \mathfrak{F}(\ell) \vdash \mathfrak{P}(\text{succ}(\ell))$ . If  $\ell$  is a **while** label then  $\text{Pre}(\ell)$  and  $\text{Post}(\ell)$  are replaced by  $\text{Inv}(\ell)$ .

This structure means that the proof is *modular*, i.e., each atomic command, function and loop is given a specification “in isolation”, without a reference to the particular context in which the specification is being used. The right-hand side of Fig. 9 shows a proof for the functions **f** and **g** which is in this form. Our approach is agnostic to the method by which the proof  $\mathfrak{P}$  is created: it can be discovered automatically (e.g., by a tool such as Abductor [9]), prepared by a proof assistant, or written manually.

**Proof assertions inlining.** The proof  $\mathfrak{P}$  assumes use of a modular specification for each function. However, our analysis needs to refer to the assertions in the local function with respect to the variables of the caller, all the way up to the top-most **work** function. We therefore define a process for lifting local assertions with respect to their global contexts.

For  $\gamma \in \text{AP}$  such that  $|\gamma| = n$  and  $\ell_1, \dots, \ell_n$  are the labels corresponding to the function calls (in the order of occurrence as in  $\gamma$ ) we define the lifted (“inlined”) proof assertion  $\mathfrak{P}(\gamma)$  as

$$\mathfrak{F}(\ell_1) * (\mathfrak{F}(\ell_2) * \dots (\mathfrak{F}(\ell_n) * \mathfrak{P}(\gamma[m])[\Omega(\ell_n)])[\Omega(\ell_{n-1})] \dots)[\Omega(\ell_1)]$$

Intuitively, the assertion  $\mathfrak{P}(\gamma)$  represents the “global” proof state at  $\gamma$  (including the framed portions) in terms of **work**’s variables.

Finally, let  $\text{pc}: \text{AP} \rightarrow \Pi_{\vee}$  represent the path constraint associated with each path delimiter. The path constraint at  $\gamma$  comprises the pure part of  $\mathfrak{P}(\gamma)$  corresponding to the assumptions from the conditionals encountered on the way from **work<sub>s</sub>** to  $\gamma$ . The path constraint may be extracted directly from the proof or computed by some other means.

## 5 Parallelization Algorithm

We now formally define our parallelization algorithm. The goal of our approach is to construct a parallelized version of the input program  $\mathbb{P}$ . In particular, our approach generates a new function **work'** in  $\mathbb{P}_{\text{par}}$  (invoking possibly transformed callees) such that

```

main $\mathbb{P}_{\text{par}}$ () {
  // initial channel creation.
  for( $C_1; C_2; b$ ) {
    // channel creation.
    fork(work', ...);
  }
  // channel finalization.
}

```

has the same behaviour as the original  $\text{main}_{\mathbb{P}}$ .

### 5.1 Resource Usage Analysis

Our approach to parallelization traverses the program by referring to a finite prefix- and  $\leftarrow$ -closed subset of path delimiters  $P \subseteq_{\text{fin}} \text{AP}$ . This subset reflects the portions of

---

**Algorithm 1** Computing locally needed resources using backward symbolic execution.

---

**function** NEEDED-LOC( $(\ell', \ell'') : \text{Label} \times \text{Label}, \Delta : \mathcal{P}(\text{SH})$ )  
 $\ell := \ell''$   
**while**  $\ell \neq \ell'$  **do**  
 $\ell := \text{pred}(\ell)$   
**if**  $\text{cmd}(\ell)$  matches **if**( $b$ ) {  $C$  } **else** {  $C'$  } **then**  
 $\Delta := \mathfrak{P}(\ell)^\Pi \wedge (\text{NEEDED-LOC}((\ell_s^t, \ell_e^t), \Delta) \cup \text{NEEDED-LOC}((\ell_s^f, \ell_e^f), \Delta))$   
**if**  $\text{cmd}(\ell)$  matches **while**( $b$ ) {  $C$  } **then**  
 $\text{Inv}(\ell)[\Omega(\ell)] * [\Delta_A] \vdash \Delta * [-]$   
 $\Delta := \mathfrak{P}(\ell)^\Pi \wedge (\text{Inv}(\ell)[\Omega(\ell)] * \Delta_A)$   
**else**  
 $\text{Post}(\ell)[\Omega(\ell)] * [\Delta_A] \vdash \Delta * [-]$   
 $\Delta := \mathfrak{P}(\ell)^\Pi \wedge (\text{Pre}(\ell)[\Omega(\ell)] * \Delta_A)$   
**return**  $\Delta$

---

the program that are path- and context-sensitive targets of parallelization. The set  $\mathsf{P}$  provides precise control over which functions the analysis should address, but how this set is chosen is not considered here. Function invocations whose successors are not in  $\mathsf{P}$  are treated as a single operation with effects defined by their specification.

The goal of resource usage analysis is to compute the maps:

$$\begin{aligned} \text{redundant} : \mathsf{P} \times \mathsf{P} &\rightarrow \mathcal{P}(\text{SH}) \\ \text{needed} : \mathsf{P} \times \mathsf{P} &\rightarrow \mathcal{P}(\text{SH}) \end{aligned}$$

For  $p, q \in \mathsf{P}$  such that  $p \prec q$ ,  $\text{redundant}(p, q)$  gives the resources that are guaranteed to not be accessed by the program between  $p$  and  $q$ . In parallelization, these are resources that can safely be transferred to other parallel threads. For  $p \prec q$ ,  $\text{needed}(p, q)$  gives the resources that might be accessed during execution from  $p$  to  $q$ . In parallelization, these are the resources that must be acquired before execution of the current thread can proceed.

The function NEEDED-LOC (Alg. 1) uses backward symbolic execution to compute needed resources between pairs of path delimiters  $(\ell', \ell'')$  in the same function block. It uses the pure parts of the sequential proof to guide the abductive inference. Since we already have function summaries and loop invariants in the sequential proof, NEEDED-LOC gives a symbolic heap sufficient to execute the fragment  $\mathbb{P}_{[\ell', \ell'']}$ .

**Lemma 2.**  $\text{NEEDED-LOC}(\ell', \ell'')$  is a sufficient precondition for  $\mathbb{P}_{[\ell', \ell'']}$ .

*Proof.* Follows from the disjunctive version of the frame rule:

$$\frac{\{P\} C \{ \bigvee_{i \in \mathcal{I}} Q_i \} \quad \Delta * \bigvee_{j \in \mathcal{J}} \Delta_j^A \vdash P * \bigvee_{k \in \mathcal{K}} \Delta_k^F}{\{ \bigvee_{j \in \mathcal{J}} (\Delta * \Delta_j^A) \} C \{ \bigvee_{i \in \mathcal{I}, k \in \mathcal{K}} (Q_i * \Delta_k^F) \}}$$

and Hoare's rule of composition. □

The function NEEDED( $\gamma_e, \gamma_s$ ) (Alg. 2) lifts NEEDED-LOC to the context-sensitive inter-procedural level. Given two assertion points represented as path delimiters  $\gamma_s$  and  $\gamma_e$ ,

---

**Algorithm 2** Computing needed resources.
 

---

```

function NEEDED( $\gamma_s: \mathbb{P}, \gamma_e: \mathbb{P}$ )
   $k := |\gamma_e|$ 
   $\Delta := \mathfrak{P}(\gamma_e)$ 
  while  $k > |\gamma_s \wedge \gamma_e| + 1$  do
     $\varrho := \Omega(\gamma_e[1..k])$ 
     $\Delta := \text{NEEDED-LOC}((\gamma_e[k]_{\uparrow}, \gamma_e[k]), \Delta[\varrho^{-1}])[\varrho]$ 
     $k := k - 1$ 
    if  $\text{cmd}(\gamma_e[k])$  is function call then
       $\Delta := \Delta[\Omega(\gamma_e[k])^{-1}]$ 
   $\varrho := \Omega(\gamma_s[1..k])$ 
   $\Delta := \text{NEEDED-LOC}((\gamma_e[k], \gamma_s[k]), \Delta[\varrho^{-1}])[\varrho]$ 
  while  $k < |\gamma_s|$  do
    if  $\text{cmd}(\gamma_s[k])$  is function call then
       $\Delta := \Delta[\Omega(\gamma_s[k])]$ 
     $k := k + 1$ 
     $\varrho := \Omega(\gamma_s[1..k])$ 
     $\Delta := \text{NEEDED-LOC}((\gamma_s[k], \gamma_s[k]_{\downarrow}), \Delta[\varrho^{-1}])[\varrho]$ 
  return  $\Delta$ 

```

NEEDED( $\gamma_e, \gamma_s$ ) works by successively pushing backwards an assertion  $\Delta$  from  $\gamma_e$  to  $\gamma_s$ . The algorithm operates in two phases. In phase A, it steps backwards from  $\gamma_e$  towards the outermost calling context in the function-invocation hierarchy. This context, represented as the longest common prefix of  $\gamma_s$  and  $\gamma_e$ , is the dominator of the two functions in which  $\gamma_s$  and  $\gamma_e$  are found in the function call graph. Phase B of the algorithm keeps stepping backwards, but proceeds inwards into the function-invocation hierarchy towards  $\gamma_s$ . Both phases of the algorithm use Alg. 1 to compute the needed resources in-between function call boundaries: in phase A we establish the needed assertions from the dominating point to  $\gamma_e$ , and in phase B from  $\gamma_s$  to the dominating point.

Since the invariants of the input proof are written in terms of the outermost calling context, comparing locally-computed specifications with these invariants requires the local specifications to be recast in terms of the outer context. In the first line of phase A we construct a variable substitution  $\varrho$  that recasts the assertion in terms of the calling context at the start of  $\gamma_e$ . The second line constructs  $\Delta[\varrho^{-1}]$ —the running state recast in terms of  $\gamma_e$ 's starting context; this is typically the context defined by the `work()` function used in a `pfor` command. NEEDED-LOC constructs a new needed state up to the start of the current block. Finally,  $\varrho$  recasts the resulting state back into the current context. When a function call is reached, we unwind the variable substitution by one call since we now have moved from the callee's context to a caller's. Operations in phase B are similar.

The results computed by NEEDED are tabulated as follows. If  $\mathfrak{P}(\gamma_s) \vdash \text{NEEDED}(\gamma_s, \gamma_e) * [-]$ , then  $\text{needed}(\gamma_s, \gamma_e) := \mathfrak{P}(\gamma_s)^{\Pi} \wedge \text{NEEDED}(\gamma_s, \gamma_e)$ ; otherwise,  $\text{needed}(\gamma_s, \gamma_e) := \mathfrak{P}(\gamma_s)$ .

**Lemma 3.** *needed( $\gamma_s, \gamma_e$ ) is a sufficient precondition to execute  $\mathbb{P}$  from  $\gamma_s$  to  $\gamma_e$ .*

In order to be able to use the `needed` map further in the algorithm we must ensure that it grows monotonically, i.e., that  $\forall \gamma, \gamma', \gamma'' \in \mathbb{P}$  such that  $\gamma \prec \gamma' \prec \gamma''$  we have  $\text{needed}(\gamma, \gamma'') \vdash \text{needed}(\gamma, \gamma') * [-]$ . If the underlying theorem prover behaves consistently



---

**Algorithm 3** Computing redundant resources
 

---

```

function REDUNDANT( $\gamma_s: \mathbb{P}, \gamma_e: \mathbb{P}$ )
   $\mathfrak{P}(\gamma_s) \vdash \text{NEEDED}(\gamma_s, \gamma_e) * [\Delta_R]$ 
  return  $\Delta_R$ 

```

---

	$\mathbf{f}_e$
$\ell_1$	$x < \mathbf{i} \wedge \mathbf{y} \mapsto y$
$\ell_2$	$(\mathbf{v} = x \wedge \mathbf{v} \geq \mathbf{i} \wedge \mathbf{x} \mapsto x)$ $\vee (\mathbf{v} = x \wedge \mathbf{v} < \mathbf{i} \wedge \mathbf{y} \mapsto y)$
$\ell_2\ell_3$	$\mathbf{v} = x \wedge \mathbf{v} \geq \mathbf{i} \wedge \mathbf{x} \mapsto x$
$\ell_2\ell_3\mathbf{g}_e$	$\mathbf{v} = x \wedge \mathbf{v} \geq \mathbf{i} \wedge \mathbf{x} \mapsto x * \mathbf{y} \mapsto y$
$\ell_2\ell_4$	$\mathbf{v} = x \wedge \mathbf{v} < \mathbf{i} \wedge \mathbf{y} \mapsto y$
$\ell_2\ell_4\mathbf{g}_e$	$\mathbf{v} = x \wedge \mathbf{v} < \mathbf{i} \wedge \mathbf{x} \mapsto x * \mathbf{y} \mapsto y$

Figure 10: redundant map with respect to  $\mathbf{f}_e$ .

with respect to failing and precision this property always holds. However, we can also make an additional check and insert assertions from the sequential proof as needed.

The redundant resource between two path delimiters is the portion of the inlined proof-state in  $\mathfrak{P}$  that is not required by the needed map. In Alg. 3 we calculate this by frame inference.

**Lemma 4.** *If  $\mathfrak{P}(\gamma_s) \vdash \text{redundant}(\gamma_s, \gamma_e) * [\Delta]$  then  $\Delta$  is a sufficient precondition to execute  $\mathbb{P}$  from  $\gamma_s$  to  $\gamma_e$ .*

Consider the functions  $\mathbf{f}$  and  $\mathbf{g}$  from §2. Fig. 9 shows the labels and sequential proof for these functions. The map obtained by NEEDED is shown in Fig. 1. Fig. 10 shows the map computed by REDUNDANT with respect to path delimiter  $\mathbf{f}_e$ . The full set of intermediate assertions are shown in two tables in Appendix C.

## 5.2 Parallelising Transformation

We now describe a parallelising transformation based on our resource-usage analysis. The construction proceeds in two phases. First, we compute an idealized resource transfer between threads. Then we inject **grant** and **wait** barriers to realise this resource transfer. The resource transfer mechanism transfers a resource from one invocation of the **work** function to another.

This parallelising transformation can be viewed as just one application of the resource-usage analysis; more optimized proof-preserving parallelising transformations are certainly possible. Our overall goal is to describe a framework for a resource-sensitive dependency-preserving analysis.

**Conditions on released and acquired resources.** In the first phase we determine resources that should be released and acquired at particular points in the parallelized program. Released resources cannot be revoked, i.e., each released resource should be included in the **redundant** map from the point of the release to the end of the **work** function—this way we know the resource will not be needed further. Acquired resources

	$l_2$	$l_2l_3$	$l_2l_3g_e$	$l_2l_4$	$l_2l_4g_e$	$f_e$
$l_1$	$\mathbf{x} \mapsto x$	$x \geq i \wedge \mathbf{x} \mapsto x$	$x \geq i \wedge \mathbf{x} \mapsto x * \mathbf{y} \mapsto y$	$x < i \wedge \mathbf{x} \mapsto x$	$x < i \wedge \mathbf{x} \mapsto x$	$(x \geq i \wedge \mathbf{x} \mapsto x * \mathbf{y} \mapsto y) \vee (x < i \wedge \mathbf{x} \mapsto x)$
$l_2$		$\mathbf{v} = x \wedge \mathbf{v} \geq i$	$\mathbf{v} = x \wedge \mathbf{v} \geq i * \mathbf{y} \mapsto y$	$\mathbf{v} = x \wedge \mathbf{v} < i$	$\mathbf{v} = x \wedge \mathbf{v} < i \wedge \mathbf{x} \mapsto x$	$(\mathbf{v} = x \wedge \mathbf{v} \geq i * \mathbf{y} \mapsto y) \vee (\mathbf{v} = x \wedge \mathbf{v} < i \wedge \mathbf{x} \mapsto x)$
$l_2l_3$			$\mathbf{v} = x \wedge \mathbf{v} \geq i * \mathbf{y} \mapsto y$			$\mathbf{v} = x \wedge \mathbf{v} \geq i * \mathbf{y} \mapsto y$
$l_2l_3g_e$						$\mathbf{v} = x \wedge \mathbf{v} \geq i$
$l_2l_4$					$\mathbf{v} = x \wedge \mathbf{v} < i * \mathbf{x} \mapsto x$	$\mathbf{v} = x \wedge \mathbf{v} < i * \mathbf{x} \mapsto x$
$l_2l_4g_e$						$\mathbf{v} = x \wedge \mathbf{v} < i$

Table 1: needed map (some entries omitted).

are held by the executing thread until released. Resources that are acquired along a sequence of path delimiters should contain what is prescribed by the **needed** map between each of the path delimiters.

We represent the result of this phase of the algorithm via the following maps:

- **resource**:  $\text{ResId} \rightarrow \mathcal{P}(\text{SH})$ , denoting resource identifiers that identify released and acquired resources;
- **released**:  $\text{P} \rightarrow \text{ResId} \times \text{Subst}$ , representing resources that are going to be released at a path delimiter together with the variable substitution applied at that point;
- **acquired**:  $\text{P} \rightarrow \text{ResId}$ , representing resources that are going to be acquired at a path delimiter.

We require the following well-formedness properties:

1.  $\forall \gamma \in \text{dom}(\text{released}) . \forall \gamma' \succ \gamma . \text{released}(\gamma) = (r, \rho) \rightarrow (\text{redundant}(\gamma, \gamma') \vdash \text{resource}(r)[\rho] * [-]);$
2.  $\forall \gamma \in \text{P} . \otimes_{r \in \text{dom}(\text{resource})} \{ \text{resource}(r) \mid \exists \gamma' \prec \gamma \wedge \text{acquired}(\gamma') = r \} \vdash \text{needed}(\gamma, \text{work}_e) * [-];$
3.  $\forall \gamma \in \text{dom}(\text{released}) . \otimes_{r \in \text{dom}(\text{resource})} \{ \text{resource}(r) \mid \exists \gamma' \prec \gamma \wedge \text{acquired}(\gamma') = r \} \vdash \text{released}(\gamma) * [-].$

The first property states we can release only resources that are not needed between the given path delimiter and any subsequent one. The second property states that the resources needed at a path delimiter must have already been acquired. The third property states that only the resources that have been previously acquired can be released.<sup>6</sup>

In general, there are many solutions satisfying properties 1–3. For instance, there is always a trivial solution that acquires  $\text{needed}(\text{work}_s, \text{work}_e)$  before the first command and releases it after the last, causing each invocation of **work** to be blocked until the preceding invocation finishes the last command. Of course, some solutions are better than others.

**Computing released and acquired maps.** Algorithm 4 constructs **released**, **acquired** and **resource** maps satisfying properties 1–3. Each iteration of the algorithm heuristically picks a needed resource, and then iteratively searches for matching redundant resource along all paths. The algorithm maintains a set  $\mathcal{C}$  of all path delimiters up to which no more resources are needed. It terminates once no unsatisfied needed resources remain (line 3).

At the start of the main loop (line 4) the algorithm picks a still-needed resource between a path delimiter in  $\mathcal{C}$  and some further path delimiter. The picking of the needed resource is governed by a heuristic function *choose*, for which we make no assumption. The *choose* function serves as a proxy for external knowledge about likely points for parallelization.

---

<sup>6</sup>We could relax the third requirement if we extended our barriers to support *renunciation* [16], the ability to release a resource without first acquiring it. Renunciation allows a resource to ‘skip’ iterations, giving limited out-of-order signalling. We believe it would be straightforward to fold such techniques into the analysis, although such extensions are outside the focus of this paper.

---

**Algorithm 4** Computing released and acquired resources.

---

```
1:  $\mathcal{N} := \text{needed}$ ;  $\mathcal{R} := \text{redundant}$ 
2:  $\mathcal{C} := \max\{\gamma \in \mathbf{P} \mid \mathcal{N}(\text{work}_s, \gamma) = \text{emp}\}$ 
3: while  $\mathcal{C} \neq \{\text{work}_e\}$  do
4:    $\Sigma_R := \mathcal{N}(\text{choose}(\{(\gamma, \gamma') \mid \gamma \in \mathcal{C}, \gamma' \in \mathbf{P}\}))^\Sigma$ 
5:    $\varrho := \bar{x} \mapsto \bar{x}'$ , where  $\bar{x}'$  fresh
6:    $\Sigma'_R := \Sigma_R[\varrho]$ 
7:    $\mathcal{C}_r := \min\{\gamma \in \mathbf{P} \mid \mathcal{R}(\gamma, \text{work}_e)^\Sigma \vdash \Sigma'_R * [-] \wedge \exists \gamma' \in \mathcal{C}. \gamma' \preceq \gamma\}$ 
8:   if  $\bigvee_{\gamma \in \mathcal{C}_r} \text{pc}(\gamma) \Leftrightarrow \text{true}$  then
9:      $r :=$  fresh resource id
10:     $\text{resource}(r) := \Sigma'_R$ 
11:    for all  $\gamma \in \mathcal{C}_r$  do
12:       $\text{released}(\gamma) := (r, \varrho)$ 
13:      for all  $\gamma'$  s.t.  $\gamma \preceq \gamma'$  do
14:         $\mathcal{R}(\gamma', \text{work}_e)^\Sigma \vdash \Sigma'_R * [\Delta]$ 
15:         $\mathcal{R}(\gamma', \text{work}_e) := \Delta$ 
16:      for all  $\gamma \in \mathcal{C}$  do
17:         $\text{acquired}(\gamma) := r$ 
18:        for all  $\gamma', \gamma''$  s.t.  $\gamma \preceq \gamma' \preceq \gamma''$  do
19:           $\mathcal{N}(\gamma', \gamma'') * [-] \vdash \Sigma'_R * [\Delta]$ 
20:           $\mathcal{N}(\gamma', \gamma'') := \Delta$ 
21:     $\mathcal{C} := \max\{\gamma' \mid \exists \gamma \in \mathcal{C}. \mathcal{N}(\gamma, \gamma') = \text{emp}\}$ 
```

---

The key step of the algorithm is performed on line 7:

$$\mathcal{C}_r := \min\{\gamma \in \mathbf{P} \mid \mathcal{R}(\gamma, \text{work}_e)^\Sigma \vdash \Sigma'_R * [-] \wedge \exists \gamma' \in \mathcal{C}. \gamma' \preceq \gamma\}$$

Here  $\mathcal{R}(\gamma, \text{work}_e)$  is the redundant resource from  $\gamma$  to the end of the `work` function and  $\Sigma'_R$  is the candidate resource that we want to acquire. The constructed set  $\mathcal{C}_r$  is a set of path delimiters along which we can satisfy the candidate needed resource. In line 8, the algorithm checks that  $\mathcal{C}_r$  covers all paths by checking the conjunction of path constraints is tautologous.

Resources stored in `needed` contain path constraints (and other conditions on local variables) embedded in the pure part of the symbolic heap. Since we can transfer resources between different path delimiters, we only take the spatial part of the resource into consideration when asking entailment questions; this is denoted by a superscript  $\Sigma$ . Moreover, since the acquired resource is being sent to a different function invocation, we substitute a fresh set of variables (line 5).

The remainder of the algorithm is devoted to constructing the new resource (line 10), and with updating `released` (lines 11–15), `acquired` (lines 16–20), and  $\mathcal{C}$  (line 21).

**Lemma 5.** *Maps resource, released and acquired computed by Algorithm 4 satisfy properties 1–3.*

Consider our running example. If `choose` picks  $(\ell_1, \ell_2)$  in the first iteration and  $(\ell_2\ell_3, \ell_2\ell_3\mathbf{g}_e)$  in the second iteration, then the end result of Alg. 4 is  $\text{resource} = \{r_1 \mapsto (\mathbf{x} \mapsto x), r_2 \mapsto (\mathbf{y} \mapsto y)\}$ ,  $\text{released} = \{\ell_2\ell_3 \mapsto (r_1, \emptyset), \ell_2\ell_4\mathbf{g}_e \mapsto (r_1, \emptyset), \ell_2\ell_3\mathbf{g}_e \mapsto (r_2, \emptyset), \ell_2\ell_4 \mapsto (r_2, \emptyset)\}$  and  $\text{acquired} = \{\ell_1 \mapsto r_1, \ell_2\ell_3\ell_5 \mapsto r_2, \ell_2\ell_4 \mapsto r_2\}$ .

**Inserting grant and wait barriers.** In this phase we transform the sequential program  $\mathbb{P}$  into a parallel program  $\mathbb{P}_{par}$  by inserting **grant** and **wait** barriers. The inserted barriers realise resource transfer defined by the maps **released** and **acquired**.

We generate the parallel function  $\mathbf{work}'(\bar{\mathbf{i}}_r^{(p)}, \bar{\mathbf{i}}_r, \mathbf{env}^{(p)}, \mathbf{env})$  in  $\mathbb{P}_{par}$  as follows:

1. To each  $r \in \mathbf{ResId}$  we assign a unique channel name  $\mathbf{i}_r$ . Denote by  $\mathbf{i}_r^{(p)}$  the corresponding channel of the previous thread in the sequence.
2. Let  $\mathbf{env}$  be an associative array that for each channel maps (escaped) local variable names to values. Let  $\mathbf{env}^{(p)}$  be such map from the previous thread in the sequence.  $\mathbf{env}$  and  $\mathbf{env}^{(p)}$  are used for *materialization*.
3. For each  $\gamma = \ell_1 \dots \ell_n \in \mathbf{dom}(\mathbf{released}) \cup \mathbf{dom}(\mathbf{acquired})$  let  $\ell_{k_1}, \dots, \ell_{k_m}$  be the labels in  $\gamma$  corresponding to function calls. Then for each  $\gamma_j := \ell_1 \dots \ell_{k_j}$  we create in  $\mathbb{P}_{par}$  an identical copy  $\mathbf{f}'$  of the function  $\mathbf{f}$  called at  $\ell_{k_j}$  and replace the call to  $\mathbf{f}$  with the call to  $\mathbf{f}'$ . Let us denote by  $\mathbf{tr}(\gamma')$  the path delimiter in  $\mathbb{P}_{par}$  corresponding to  $\gamma'$  after this transformation has been applied for all  $\gamma \in \mathbf{dom}(\mathbf{released}) \cup \mathbf{dom}(\mathbf{acquired})$ .
4. For each  $\gamma \in \mathbf{dom}(\mathbf{acquired})$  such that  $\mathbf{acquired}(\gamma) = r$  we insert a wait barrier  $\mathbf{wait}(\mathbf{i}_r^{(p)})$  between path delimiters  $\mathbf{tr}(\mathbf{pred}(\gamma))$  and  $\mathbf{tr}(\gamma)$ .
5. For each  $\gamma \in \mathbf{dom}(\mathbf{released})$  such that  $\mathbf{released}(\gamma) = (r, \_)$ , between path delimiters  $\mathbf{tr}(\mathbf{pred}(\gamma))$  and  $\mathbf{tr}(\gamma)$  we insert a sequence of assignments of the form  $\mathbf{env}(\mathbf{i}_r)[\mathbf{y}] := \mathbf{y}$  for every local variable  $\mathbf{y}$ , followed by a grant barrier  $\mathbf{grant}(\mathbf{i}_r)$ .

Each invocation of **work** creates a fresh set of local variables that are bound to the scope of the function. However, some resources must be expressed in terms of function-local variables. Parallelization must take account of this. If the structure of a resource depends on local variables from a previous invocation, this must be encoded explicitly by *materialising* the variables of the previous invocation.

The main function  $\mathbf{main}_{\mathbb{P}_{par}}$  in  $\mathbb{P}_{par}$  first creates the set of “dummy” channels; then in the **while** loop repeatedly creates a set of new channels for the current iteration, forks a new thread with **work'** taking the channels from the previous iteration as  $\bar{\mathbf{i}}_r^{(p)}$  and from the current iteration as  $\bar{\mathbf{i}}_r$ , and at the end of the loop body assigns the new channels to the previous channels; and, after the **while** loop completes waits on the channels in the last set.

We generate the parallel proof  $\mathfrak{P}_{par}$  from the sequential proof  $\mathfrak{P}$  using the following specifications from [16]:

$$\begin{array}{ccc} \{\mathbf{emp}\} & \mathbf{i} := \mathbf{newchan}() & \{\mathbf{req}(\mathbf{i}, R) * \mathbf{fut}(\mathbf{i}, R)\} \\ \{\mathbf{req}(\mathbf{i}, R) * R\} & \mathbf{grant}(\mathbf{i}) & \{\mathbf{emp}\} \\ \{\mathbf{fut}(\mathbf{i}, R)\} & \mathbf{wait}(\mathbf{i}) & \{R\} \end{array}$$

Each variable  $R$  in  $\mathfrak{P}_{par}$  associated with channel  $\mathbf{i}_r$  is instantiated with the corresponding resource  $\mathbf{resource}(r)$ . The predicates **req** and **fut** track the ownership of the input and output ends of the channel. To reason about threads, we use the standard separation logic rules for fork-join disjoint concurrency (e.g., as in [18]).

**Theorem 6.**  $\mathfrak{P}_{par}$  is a proof of the parallel program  $\mathbb{P}_{par}$ , and defines the same specification for  $\mathbf{main}_{\mathbb{P}_{par}}$  as  $\mathfrak{P}$  does for  $\mathbf{main}_{\mathbb{P}}$ .

**Loop-splitting.** The approach presented so far treats a loop as a single command with a specification derived from its invariant. Acquiring or releasing resources within a loop is subtle, as it changes the sequential loop invariant. It is not clear how to handle this in full generality, so we take a pragmatic approach that performs heuristic loop-splitting.

The example in §3 uses two channels to transfer the segment of the list traversed after the first and the second `while` loop, respectively. The resource released via channel `i1` in Fig. 7 is  $\text{hd} \mapsto h * \text{lseg}(h, \text{xpr})$ . In the following iteration, the needed resource for the whole loop is  $\text{hd} \mapsto h * \text{lseg}(h, \mathbf{x})$ . If we try to match released against needed, the entailment  $\mathcal{R}(\gamma, \text{work}_e)^\Sigma \vdash \Sigma'_R * [-]$  in Algorithm 4 will fail. This is because the value of  $\mathbf{x}$  is unknown at the start of the loop, meaning we cannot establish whether the released resource will cover the needed resource.

One way to resolve this would be to acquire the entire list before the first loop, but this would result in a very poor parallelization. Instead, we modify the structure of the loop to expose the point at which the second list segment becomes necessary:

1. We split the spatial portion of the resource needed by the whole loop into a “dead” (already traversed) and a “live” (still to be traversed) part. In our example,  $\text{hd} \mapsto h * \text{lseg}(h, \mathbf{x})$  would be the “dead” and  $\text{lseg}(\mathbf{x}, \text{nil})$  the live part. This kind of splitting is specific to some classes of programs, e.g., linked list programs that do not traverse a node twice.
2. We match the resource against the “dead” part of the loop invariant and infer the condition under which the two resources are the same. In our example, the entailment between the two resources holds if  $\mathbf{x} = \text{xpr}$ . This condition can be inferred by asking a bi-abduction question  $\mathcal{R}(\gamma, \text{work}_e)^\Sigma \wedge [c] \vdash \Sigma'_R * [-]$  with the pure abducted fact  $c$ .

Now we can syntactically split the loop against the inferred condition  $c = (\mathbf{x} = \text{xpr})$  and obtain a transformed version that ensures that after entering the true branch of the `if` statement the condition  $c$  holds. The transformation of our example is shown in Fig 11.

Formally, we define splitting of a command (comprising possibly multiple loops) against a condition  $c$  as follows (to simplify, we assume here that every command ends with `skip`):

```

function SPLIT( $C$ : Cmd,  $c$ :  $\Pi_\vee$ )
  match  $C$  with
    | while( $b$ ) {  $C'$  };  $C'' \rightarrow$ 
      while( $b \wedge \neg c'$ ) {  $C'$  };
      if( $b \wedge c'$ ) {  $C'$ ; while( $b$ ) {  $C'$  };  $C''$  }
      else { SPLIT( $C''$ ,  $c$ ) }
    |  $C'$ ;  $C'' \rightarrow C'$ ; SPLIT( $C''$ ,  $c$ )
    | skip  $\rightarrow$  skip

```

The condition  $c'$  is obtained from  $c$  by replacing a reference to every primed variable  $y'$  by a reference to  $\text{env}^{(p)}(i_r^{(p)})[y']$ , where  $i_r$  is the channel name associated to the resource. It is not difficult to see that the accompanying proof of  $C$  can be split in a proof preserving way against  $c$ . This transformation can either be applied to  $\mathbb{P}$  between the resource-usage analysis and parallelization, or embedded within Alg. 4.

```

...
xpr:=envp(i1) ["x"];
while(x!=nil&&i!=n&&x!=xpr){
  sum+=x.val;
  i++;
  x:=x.nxt;
}
if(x!=nil&&i!=n&&x==xpr){
  sum+=x.val;
  i++;
  x:=x.nxt;
  while(x!=nil&&i!=n){
    sum+=x.val;
    i++;
    x:=x.nxt;
  }
  // remainder skipped.
...
else{
...

```

Figure 11: Loop-splitting for the `sum_head` example.

### 5.3 Implementation

We have validated our parallelization algorithm by crafting a prototype implementation on top of the existing separation logic tool, `coreStar` [7]. While our implementation is not intended to provide full end-to-end automated translation, it is capable of validating the algorithms on the examples given in the paper, and automatically answering the underlying theorem proving queries.

Our parallelization algorithm does not assume a shape invariant generator, except possibly to help construct the sequential proof. Soundness is independent of the “cleanliness” of the invariants (the analysis will always give a correct result, in the worst case defaulting to sequential behaviour). Our examples in `coreStar` [7] have been validated using automatically-generated invariants. Other efforts [9] indicate that bi-abduction works well with automatically-generated invariants produced by shape analysis, even over very large code bases.

## 6 Behaviour Preservation

A distinctive property of our parallelization analysis is that it enforces sequential data-dependencies in the parallelized program even if the safety proof does not explicitly reason about such dependencies. The result is that our analysis preserves the sequential behaviour of the program: any behaviour exhibited by the parallelized program is also a behaviour that could have occurred in the original sequential program. However, there are important caveats relating to *termination* and *allocation*.

*Termination.* If the original sequential program does not terminate, our analysis may

introduce new behaviours simply by virtue of running segments of the program that would be unreachable under a sequential schedule. To see this, suppose we have a `pfor` such that the first iteration of the loop will never terminate. Sequentially, the second iteration of the loop will never execute. However, our parallelization analysis will execute all iterations of the loop in parallel. This permits witnessing behaviours from the second (and subsequent) iterations. These behaviours were latent in the original program, and become visible only as a result of parallelization.

*Allocation and disposal.* If the program both allocates and disposes memory, the parallelized program may exhibit aliasing that could not occur in the original program. To see this, consider the following sequential program:

```
x=alloc(); y=alloc(); dispose(x).
```

For simplicity, we have avoided re-structuring the program to use data parallelism via `pfor`—this example could easily be encoded as such, however. Parallelization might give us the following program:

```
(x=alloc(); grant(wx); y=alloc()) || (wait(wx); dispose(x))
```

This parallelized version of the program is race-free and obeys the required sequential ordering on data. Depending upon the implementation of the underlying memory allocator, however, `x` and `y` may be aliased if the `dispose` operation was interleaved between the two allocations. Such aliasing could not happen in the original non-parallelized version.

Either kind of new behaviour might result in further new behaviours—for example, we might have an if-statement conditional on `x==y` in the second example above. These caveats are common to our analysis and others based on separation logic—for example, see the similar discussion in [12].

**Proving behaviour preservation.** We now sketch a behaviour preservation result (a detailed proof is given in Appendix). The theorem defines preservation in terms of an interleaved operational semantics for a core language, similar to the one described in §4.1, additionally equipped with threads, and operations on channels (such as `grant` and `wait`). Data-races are interpreted as faults in this semantics. We prove our result for a sequential thread  $t$  possibly executing concurrently with other sequential threads—this degenerates into the purely sequential case when there are no other runnable threads. Because  $t$  is sequential, it does not call `fork`; we also assume that it is guaranteed to terminate, and that it never disposes memory based on the caveats discussed above.

**Theorem 7.** *Let  $t$  be a sequential thread, and let  $t_{\text{par}}$  be a corresponding parallelized version, equipped with thread creation operations, and synchronization actions that enforce sequential dependencies among the child threads it creates. Let  $\mathcal{K}_{\text{par}}$  be a terminating non-faulting (i.e., data-race free) trace of an execution of  $t_{\text{par}}$ . There exists a corresponding trace  $\mathcal{K}$  derived by substituting corresponding operations in  $t$  for  $t_{\text{par}}$  such that: (1)  $\mathcal{K}$  and  $\mathcal{K}_{\text{par}}$  begin in the same state; (2) for every thread  $t' \neq t$  in  $\mathcal{K}$ ,  $\mathcal{K}$  and  $\mathcal{K}_{\text{par}}$  exhibit identical thread-local behaviour; and (3) the terminal state of thread  $t$  in  $\mathcal{K}_{\text{par}}$  is a substate of the corresponding state in  $\mathcal{K}$ .*



*Proof sketch.* We show that, under the assumption that forked child threads never wait for channels granted by their parent or later-forked child threads, any terminating non-faulting trace  $\mathcal{K}_{\text{par}}$  can be reordered into a sequentialized trace  $\mathcal{K}_{\text{seq}}$  with the same thread-local behaviour. By *sequentialized*, we mean that forked children must execute to completion before their parents can be scheduled.

We establish a simulation invariant between executions of  $t_{\text{par}}$  decorated with arbitrary calls to **newch**, **wait**, **grant** and **fork**, and executions of  $t_{\text{par}}^{\downarrow}$  identical to  $t_{\text{par}}$  except with all such constructs erased. The invariant establishes that every non-faulting sequentialized trace  $\mathcal{K}_{\text{seq}}$  of  $t_{\text{par}}$  simulates some trace  $\mathcal{K}$  of  $t_{\text{par}}^{\downarrow}$ .

We tie this result to our analysis by observing that parallelization only inserts the four constructs **newch**, **wait**, **grant** and **fork** (leaving aside issues of loop-splitting and materialization). As a result,  $t_{\text{par}}^{\downarrow} = t$ . We also show that the way we insert signals ensures that they are ordered with respect to thread creation, as assumed in the previous step.

This establishes a behaviour-preservation result for non-faulting traces, but parallelization might introduce faults. However, parallelized programs are verified using separation logic, meaning we can assume that they do not fault when executed in a state satisfying the precondition. This completes the proof.  $\square$

One attractive property is that this proof does not place any explicit requirements on the positioning of barriers—it is sufficient that we can provide a separation-logic proof for  $t_{\text{par}}$ . The caveat on memory disposition manifests in the way we establish our simulation invariant, which necessarily assumes newly allocated data is always “fresh” and thus does not alias with any accessible heap-allocated structure. The caveat on termination is necessary to ensure we can suitably reorder forked threads to yield a sequentialized trace.

In addition to inserting barriers, our analysis mutates the program by materialising thread-local variables and splitting loops. Both of these can be performed as mutations on the initial sequential program, and neither affect visible behaviour. Loop-splitting is straightforwardly semantics-preserving, while materialized variables are only used to control barrier calls.

Theorem 7 guarantees that parallelisation does not introduce deadlocks; otherwise the simulation relation would not exist. The ordering on barriers ensures that termination in the sequential program is preserved in the resulting parallel program.

## 7 Refining the Predicate Domain

The structure of the sequential proof affects the success of parallelization in two ways. First, the loop invariants may allow the analysis to verify a parallelized program when it would otherwise have failed. Second, the choice of predicate domain controls how much information about resource-usage is available to the analysis. Intuitively, enriching the domain may permit a finer-grained splitting of resources, allowing redundant resources to be identified earlier or missing resources later.

To see how the choice of abstract domain influences precision and effectiveness of the analysis, consider the example discussed in §3; there, we parallelized **sum\_head** using the list segment predicate **lseg**. An alternative predicate is **lsegni**( $h, t, n, i$ ), which extends **lseg** with a parameter  $n \in \mathbb{Z}$  recording the *length* of the list segment, and with a parameter  $i \in (0, 1]$  recording the *permission* on the list segment. If  $i = 1$  the thread has read-write

access; otherwise it has read access only. We define `lsegni` as the least predicate satisfying the following equation:

$$\text{lsegni}(x, t, n, i) \triangleq (x = t \wedge n = 0 \wedge \text{emp}) \vee (\exists v, y. x.\text{val} \xrightarrow{i} v * x.\text{nxt} \xrightarrow{i} y * \text{lsegni}(y, t, n-1, i))$$

The equivalence  $\text{lseg}(h, t) \iff \exists n. \text{lsegni}(h, t, n, 1)$  would allow the analysis to exploit this finer-grained predicate even if the sequential proof is written using the coarse-grained `lseg` predicate.

Function `sum_head` only writes to the list after traversing `n` nodes. Consequently, it needs only non-exclusive, read-only access to the first `n` nodes. (Or the entire list, if the list is shorter than `n` nodes long.) We can extend our resource usage analysis to deal with this richer predicate such that it identifies the following as a sufficient precondition for the *entire* `sum_head` function:

$$\text{hd} \xrightarrow{j} h * \left( (\exists t. \text{lsegni}(h, t, n, i) * \text{lseg}(t, \text{nil})) \vee (\exists n'. \text{lsegni}(h, \text{nil}, n', i) \wedge n' < n) \right)$$

Without the `lsegni` predicate, this precondition could not be expressed. The availability of `lsegni` would also allow the analysis to discover that this resource is redundant at the start of `sum_head`:

$$\exists h, i, j. \text{hd} \xrightarrow{j} h * (\text{lsegni}(h, t, n, i) \vee (\exists n'. \text{lsegni}(h, \text{nil}, n', i) \wedge n' < n))$$

When `sum_head(n)` is called in a `parallel-for`, the subsequent call to `sum_head` can *immediately* read from the first `n` nodes of the list. Applying our analysis using the `lsegni` predicate yields the the two-thread parallelization shown in Fig. 12. `i1` signals that the first `n` nodes can be read by the subsequent iteration of `sum_head`. Similar to the transformation discussed in §3, `grant` calls can be inserted on `i2` to signal that the subsequent iteration can write to the first `n` nodes; and on `i3` to signal that the entire list can be written to. Because this parallelization calls `grant()` earlier than the parallelization discussed earlier, it can consequently extract more parallelism from the original sequential program.

As this discussion reveals, our analysis is generic in the choice of abstract domain; any separation logic predicate could be used in place of `lseg`, for example. However, the success of *automated* parallelization is highly dependent on the power of the entailment prover in the chosen domain. The `lseg` domain is one of the best-developed in separation logic, and consequently automated parallelization is feasible using tools such as `coreStar` [7]. Other domains (such as trees) are far less developed.

## 8 Related Work

**Resource-usage inference by abduction.** We have defined an interprocedural, control-flow-sensitive analysis capable of determining the resource that will (and will not) be accessed between particular points in the program. At its core, our analysis uses abductive reasoning [9] to discover *redundancies*—that is, state used earlier in the program that will not be accessed subsequent to the current program point. Using abduction in this

```

sum_head1(n){
  *npr = n;
  grant(i1);
  i = 1;
  sum = 0;
  x = *hd;
  while(x!=nil && i!=n){
    sum += x.val;
    i++;
    x = x.nxt;
  }
  *xpr = x;
  grant(i2);
  while(x!=nil){
    x.val = 0;
    x = x.nxt;
  }
  grant(i3);
}

sum_head2(n){
  i = 1;
  sum = 0;
  wait(i1);
  x = *hd;
  while(x!=nil && i!=n){
    if(n==*npr) wait(i2);
    sum += x.val;
    i++;
    x = x.nxt;
  }
  wait(i2);
  while(x!=nil){
    if(x==*xpr) wait(i3);
    x.val = 0;
    x = x.nxt;
  }
}

```

Figure 12: Improved parallelization of `sum_head`.

way was first proposed in [14], where it is used to discover memory leaks, albeit without conditionals, procedures, loops, or code specialization.

In [10], abduction is used to infer resource invariants for synchronization, using a process of counterexample-driven refinement. Our approach similarly infers resource invariants, but using a very different technique: invariants are derived from a sequential proof, and we also infer synchronization points and specialise the program to reveal synchronization opportunities.

**Behaviour-preserving parallelization.** We expect our resource-usage analysis can be used in other synchronization-related optimizations, but in this paper, we have used it as the basis for a parallelising transformation. This transformation is in the style of *deterministic parallelism* [4, 5, 3, 8]—although our approach does not, in fact, require determinacy of the source program. In this vein, our transformation ensures that every behaviour of the parallelized program is a behaviour of the source sequential program (modulo the caveats about allocation and termination discussed in §6).

Previous approaches to deterministic parallelism operate without the benefit of a high-level specification. This places a substantial burden on the analysis and runtime to safely extract information on resource usage and transfer—information that is readily available in a proof. As a result, these analyses tend to be much more conservative in their treatment of mutable data. Our proof-based technique gives us a general approach to splitting mutable resources; for example, by allowing the analysis to perform ad-hoc list splitting, as we do with `sum_head`.

Our approach transforms a sequential `for`-loop by running all the iterations in parallel and signalling between them. This idea has been proposed previously, for example in numerical computation [32]. The novelty in our approach lies in inferring synchronization

over (portions of) complex mutable data-structures. Alternatively, we could have used a more irregular concurrency annotation, for example safe futures [26], or an unordered `parallel-for`, as in the Galois system [29]. In the former case, our resource-usage analysis would be mostly unchanged, but our parallelized program would construct a set of syntactically-distinct threads, rather than a pipeline of identical threads. Removing ordering between iterations, as in the latter case, would mean replacing ordered `grant-wait` pairs with conventional locks, and would introduce an obligation to show that locks were always acquired together, as a set.

**Proof-driven parallelization.** A insight central to our approach is that a separation logic proof expresses data dependencies for *parts* of a program, as well as for the whole program. These internal dependencies can be used to inject safe parallelism. This insight is due to [30, 12] and [23], both of which propose parallelization analyses based on separation logic. The analyses proposed in these papers are much more conservative than ours, in that they discover independence which already exists between commands of the program. They do not insert synchronization constructs, and consequently cannot enforce sequential dependencies among concurrent computations that share and modify state. Indeed, [30] does not consider any program transformations, since the goal of that work is to identify memory separation of different commands, while [23] expresses optimizations as reordering rewrites on proof trees.

Bell *et. al* [1] construct a proof of an *already-transformed* multithreaded program parallelized by the DSWP transformation [28]. This approach assumes a specific pattern of (linear) dependencies in the while-loop consistent with DSWP, a specific pattern of sequential proof, and a fixed number of threads. In our `sum_head` example, the outermost (parallelising) loop contains two successive inner loops, while the example in Fig. 2 illustrates how the technique can deal with interprocedural and control-flow sensitive dependencies. In both cases, the resulting parallelization is specialized to inject synchronization primitives to enforce sequential dependencies. We believe examples like these do not fall within the scope of either DSWP or the proof techniques supported by [1].

Outside separation logic, Deshmukh *et. al* [13] propose an analysis which augments a sequential library with synchronization. This approach takes as input a sequential proof expressing the correctness criteria for the library, and generates synchronization ensuring this proof is preserved if methods run concurrently. A basic assumption is that the sequential proof represents all the properties that must be preserved. In contrast, we also preserve sequential order on access to resources. Consequently, Deshmukh *et al.* permit parallelizations that we would prohibit, and can introduce new behaviours into the parallelized program. Another difference is that [13] derives a linearizable implementation given a sequential specification in the form of input-output assertions; because they do not consider *specialization* of multiple instances of the library running concurrently, it is unclear how their approach would deal with transformations of the kind we use for `sum_head`.

**Separation logic and concurrency.** Separation logic is essential to our approach. It allows us to localise the behaviour of a program fragment to precisely the resources it accesses. Our proofs are written using concurrent separation logic [27]. CSL has been extended to deal with dynamically-allocated locks [18, 22, 24], re-entrant locks [19], and

primitive channels [21, 1, 33, 25]. Sequential tools for separation logic have achieved impressive scalability—for example [9] has verified a large proportion of the Linux kernel. Our work can be seen as an attempt to leverage the success of such sequential tools. Our experiments are built on `coreStar` [7, 15], a language-independent proof tool for separation logic.

The parallelization phase of our analysis makes use of the specifications for parallelization barriers proposed in [16]. That paper defined high-level specifications representing the abstract behaviour of barriers, and verified those specifications against the barriers’ low-level implementations. However, it assumed that barriers were already placed in the program, and made no attempt to infer barrier positions. In contrast, we assume the high-level specification, and define an analysis to insert barriers. The semantics of barriers used in that paper and here was initially proposed in [26].

## Acknowledgements

This work was supported by the Gates trust, by EPSRC grant EP/H010815/1, and by NSF grant CCF-0811631. Thanks to Dino Distefano, Matthew Parkinson, Mohammad Raza, John Wickerson and the anonymous reviewers for comments and suggestions.

## References

- [1] Christian J. Bell, Andrew Appel, and David Walker. Concurrent Separation Logic for Pipelined Parallelization. In *SAS*, pages 151–166, 2009.
- [2] Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *FMCO*, pages 115–137, 2005.
- [3] Tom Bergan, Owen Anderson, Joseph Devietti, Luis Ceze, and Dan Grossman. CoreDet: A Compiler and Runtime System for Deterministic Multithreaded Execution. *SIGPLAN Not.*, 45(3):53–64, 2010.
- [4] Emery D. Berger, Ting Yang, Tongping Liu, and Gene Novark. Grace: Safe multithreaded programming for C/C++. In *OOPSLA*, pages 81–96, 2010.
- [5] Robert L. Bocchino, Jr., Vikram S. Adve, Danny Dig, Sarita V. Adve, Stephen Heumann, Rakesh Komuravelli, Jeffrey Overbey, Patrick Simmons, Hyojin Sung, and Mohsen Vakilian. A Type and Effect System for Deterministic Parallel Java. In *OOPSLA*, pages 91–116, 2009.
- [6] Richard Bornat, Cristiano Calcagno, Peter O’Hearn, and Matthew Parkinson. Permission Accounting in Separation Logic. In *POPL*, pages 259–270, 2005.
- [7] Matko Botinčan, Dino Distefano, Mike Dodds, Radu Griore, Naudžiūnienė, and Matthew Parkinson. `coreStar`: The Core of `jStar`. In *Boogie*, pages 65–77, 2011.
- [8] Jacob Burnim and Koushik Sen. Asserting and Checking Determinism for Multithreaded Programs. *Commun. ACM*, 53:97–105, June 2010.

- [9] Cristiano Calcagno, Dino Distefano, Peter O’Hearn, and Hongseok Yang. Compositional Shape Analysis by Means of Bi-Abduction. In *POPL*, pages 289–300, 2009.
- [10] Cristiano Calcagno, Dino Distefano, and Viktor Vafeiadis. Bi-abductive Resource Invariant Synthesis. In *APLAS*, pages 259–274, 2009.
- [11] Cristiano Calcagno, Peter W. O’Hearn, and Hongseok Yang. Local Action and Abstract Separation Logic. In *LICS*, pages 366–378, 2007.
- [12] Byron Cook, Stephen Magill, Mohammad Raza, Jiri Simsa, and Satnam Singh. Making Fast Hardware with Separation Logic, 2010. Unpublished, <http://cs.cmu.edu/~smagill/papers/fast-hardware.pdf>.
- [13] Jyotirmoy V. Deshmukh, G. Ramalingam, Venkatesh Prasad Ranganath, and Kapil Vaswani. Logical Concurrency Control from Sequential Proofs. In *ESOP*, pages 226–245, 2010.
- [14] Dino Distefano and Ivana Filipović. Memory Leaks Detection in Java by Bi-abductive Inference. In *FASE*, pages 278–292, 2010.
- [15] Dino Distefano and Matthew J. Parkinson J. jStar: Towards Practical Verification for Java. In *OOPSLA*, pages 213–226, 2008.
- [16] Mike Dodds, Suresh Jagannathan, and Matthew J Parkinson. Modular Reasoning for Deterministic Parallelism. In *POPL*, pages 259–270, 2011.
- [17] Tayfun Elmas, Shaz Qadeer, and Serdar Tasiran. A Calculus of Atomic Actions. In *POPL*, pages 2–15, 2009.
- [18] Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzky, and Mooly Sagiv. Local Reasoning for Storable Locks and Threads. In *APLAS*, pages 19–37, 2007.
- [19] Christian Haack, Marieke Huisman, and Clément Hurlin. Reasoning about Java’s Reentrant Locks. In *APLAS*, pages 171–187, 2008.
- [20] Tim Harris, Jim Larus, and Ravi Rajwar. *Transactional Memory, 2nd edition*. Morgan-Claypool, 2010.
- [21] C. A. R. Hoare and Peter W. O’Hearn. Separation Logic Semantics for Communicating Processes. *ENTCS*, 212:3–25, 2008.
- [22] Aquinas Hobor, Andrew W. Appel, and Francesco Zappa Nardelli. Oracle semantics for concurrent separation logic. In *ESOP*, 2008.
- [23] Clément Hurlin. Automatic Parallelization and Optimization of Programs by Proof Rewriting. In *SAS*, pages 52–68, 2009.
- [24] Bart Jacobs and Frank Piessens. Modular full functional specification and verification of lock-free data structures. Technical Report CW 551, Katholieke Universiteit Leuven, Dept. of Computer Science, 2009.

- [25] K. R. M. Leino, P. Müller, and J. Smans. Deadlock-free Channels and Locks. In *ESOP*, pages 407–426, 2010.
- [26] A. Navabi, X. Zhang, and S. Jagannathan. Quasi-static Scheduling for Safe Futures. In *PPoPP*, pages 23–32. ACM, 2008.
- [27] Peter W. O’Hearn. Resources, Concurrency and Local Reasoning. *TCS*, 375:271–307, 2007.
- [28] Guilherme Ottoni, Ram Rangan, Adam Stoler, and David I. August. Automatic Thread Extraction with Decoupled Software Pipelining. In *MICRO*, pages 105–118, 2005.
- [29] Keshav Pingali, Donald Nguyen, Milind Kulkarni, Martin Burtscher, M. Amber Hasaan, Rashid Kaleem, Tsung-Hsien Lee, Andrew Lenharth, Roman Manevich, Mario Méndez-Lojo, Dimitrios Proutzos, and Xin Sui. The Tao of Parallelism in Algorithms. In *PLDI*, pages 12–25, 2011.
- [30] Mohammad Raza, Cristiano Calcagno, and Philippa Gardner. Automatic Parallelization with Separation Logic. In *ESOP*, pages 348–362, 2009.
- [31] John C. Reynolds. Separation Logic: A Logic for Shared Mutable Data Structures. In *LICS*, pages 55–74, 2002.
- [32] Peiyi Tang, Peiyi Tang, John N. Zigman, and John N. Zigman. Reducing Data Communication Overhead for DOACROSS Loop Nests. In *International Conference on Supercomputing*, pages 44–53, 1993.
- [33] Jules Villard, Étienne Lozes, and Cristiano Calcagno. Tracking Heaps That Hop with Heap-Hop. In *TACAS*, pages 275–279, 2010.

## A Semantics

In this section, we define an annotated operational semantics for our core programming language. Our semantics makes several simplifying assumptions: (1) We forbid disposal, and so can assume that allocation always gives fresh locations. This restriction is introduced because full memory allocation (with memory disposition and reuse) may allow parallelisation to introduce new behaviours (see §6); (2) to simplify the presentation, our semantics does not include function definitions or calls.

The operational semantics has two levels, a labelled thread-local semantics ( $\rightsquigarrow$ ), and a global semantics ( $\Vdash$ ) defined in terms of the local semantics. Every transition in the global semantics corresponds to a transition for some thread in the local semantics. Our semantics is *annotated* in the sense that it keeps track of the resources that are associated with threads and channels. These distinctions would not be present in the operational semantics of the real machine.

**Thread-local semantics.** The semantics assumes the following basic sets: **Prog**, programs; **Tid**, thread identifiers, ordered by  $<$ ; **Cid**, channel identifiers; **Var**, variable names; **Loc**, heap locations; **Val**, primitive values (which include locations). We assume that the sets **Tid**, **Cid** and **Loc** are infinite, and that  $\text{Tid} \uplus \text{Cid} \uplus \text{Loc} \subseteq \text{Val}$ .

The thread-local semantics, given in Fig. 13, is defined by a labelled transition relation  $\rightsquigarrow \in \mathcal{P}(\text{TState} \times \text{TState} \times \text{Label})$ , with **TState** defined as follows:

$$\begin{aligned} \sigma_s &\in \text{Stack} : \text{Var} \rightarrow \text{Val} \\ \sigma_h &\in \text{Heap} : \text{Loc} \rightarrow \text{Val} \\ \sigma &\in \text{LState} : \text{Stack} \times \text{Heap} \\ (\omega_w, \omega_g) &\in \text{CState} : \mathcal{P}(\text{Cid}) \times \mathcal{P}(\text{Cid}) \\ \text{TState} &: \text{Prog} \times \text{LState} \times \text{CState} \times \mathcal{P}(\text{Tid}) \end{aligned}$$

The semantics associates each thread with a resource which can be safely manipulated by that thread. Globally-visible events, meaning forking, creating channels, granting and waiting, and heap allocation are modelled by labels on the thread-local transition relation. These labels are used to ensure that global events are propagated to all threads. The set of labels, **Label**, is defined as follows:

$$\begin{aligned} \text{Label} &\triangleq (\{\mathbf{fork}\} \rightarrow \text{Tid} \times \text{Prog} \times \text{LState} \times \text{CState}) \\ &\uplus (\{\mathbf{newch}\} \rightarrow \text{Cid}) \\ &\uplus (\{\mathbf{wait}, \mathbf{grant}\} \rightarrow \text{Cid} \times \text{Heap}) \\ &\uplus (\{\mathbf{alloc}\} \rightarrow \text{Loc}) \end{aligned}$$

The rules use a join operator on heaps  $\oplus : \text{Heap} \times \text{Heap}$  by union of functions, defined only if the two heaps' domains are disjoint. The join operator  $\oplus : \text{LState} \times \text{LState} \rightarrow \text{LState}$  is defined as identity on stacks and join on heaps:

$$\sigma \oplus \sigma' \triangleq (\sigma_s, \sigma_h \oplus \sigma'_h) \quad \text{if } \sigma_s = \sigma'_s$$

A *local state* in **LState** consists of a *stack* mapping variables to values and a *heap* mapping locations to values. Note that we do not support address arithmetic. A *channel state* in **CState** consists of a pair of sets, respectively recording the channels the thread can **wait** for and **grant** on. A *thread state* in **TState** consists of a program (i.e. command), a local state, a channel state, and a set of thread identifiers of child threads.

The semantics assumes a set of primitive commands  $\text{Prim} : \mathcal{P}(\text{LState} \times \text{LState})$ . We assume for any command  $c \in \text{Prim}$  that:

- For any states  $\sigma_1$  and  $\sigma_2$  such that  $(\sigma_1 \oplus \sigma_2, \sigma') \in c$ , if there exists a transition  $(\sigma_1, \sigma'_1) \in c$  then  $\sigma'_1 \oplus \sigma_2 = \sigma'$ .
- For any transition  $(\sigma, \sigma') \in c$  and state  $\sigma_2$ , if  $\sigma \oplus \sigma_2$  is well-defined, then there exists a transition  $(\sigma \oplus \sigma_2, \sigma'') \in c$  and  $\sigma'' = \sigma' \oplus \sigma_2$ .
- For any states  $\sigma_1$  and  $\sigma_2$ , if there exists no  $\sigma'$  such that  $(\sigma_1 \oplus \sigma_2, \sigma') \in c$ , then  $(\sigma_1, \sigma') \notin c$ .



$$\begin{array}{c}
\frac{(\sigma, \sigma') \in c}{(c, \sigma, \omega, \gamma) \rightsquigarrow (\mathbf{skip}, \sigma', \omega, \gamma)} \quad \frac{\neg \exists \sigma'. (\sigma, \sigma') \in c}{(c, \sigma, \omega, \gamma) \rightsquigarrow \mathbf{abort}} \\
\\
\frac{\sigma, \omega \not\models P * \mathbf{true}}{(\mathbf{fork}_{[P]} C, \sigma, \omega, \gamma) \rightsquigarrow \mathbf{abort}} \quad \frac{\sigma', \omega' \models P \quad \sigma = \sigma' \oplus \sigma'' \quad \omega = \omega' \uplus \omega'' \quad t \text{ fresh}}{(\mathbf{fork}_{[P]} C, \sigma, \omega, \gamma) \xrightarrow{\mathbf{fork}(t, C, \sigma', \omega')} (\mathbf{skip}, \sigma'', \omega'', \gamma \uplus \{t\})} \\
\\
\frac{\mathbf{newch} \text{ fresh}}{(x := \mathbf{newch}, \sigma, \omega, \gamma) \xrightarrow{\mathbf{newch}(s)} (\mathbf{skip}, (\sigma_s[x \mapsto s], \sigma_h), (\omega_w \uplus \{s\}, \omega_g \uplus \{s\}), \gamma)} \\
\\
\frac{\sigma, (\emptyset, \emptyset) \not\models P * \mathbf{true}}{(\mathbf{grant}_{[P]} E, \sigma, \omega, \gamma) \rightsquigarrow \mathbf{abort}} \quad \frac{[E]_{\sigma_s} = s \quad s \in \omega_g \quad \sigma = \sigma' \oplus \sigma'' \quad \sigma', (\emptyset, \emptyset) \models P}{(\mathbf{grant}_{[P]} E, \sigma, \omega, \gamma) \xrightarrow{\mathbf{grant}(s, \sigma')} (\mathbf{skip}, \sigma'', (\omega_w, \omega_g \setminus \{s\}), \gamma)} \\
\\
\frac{[E]_{\sigma_s} = s \quad s \in \omega_w}{(\mathbf{wait} E, \sigma, \omega, \gamma) \xrightarrow{\mathbf{wait}(s, \sigma')} (\mathbf{skip}, \sigma \oplus \sigma', (\omega_w \setminus \{s\}, \omega_g), \gamma)} \\
\\
\frac{(C, \sigma, \omega, \gamma) \rightsquigarrow (C', \sigma', \omega', \gamma')}{(C; C'', \sigma, \omega, \gamma) \rightsquigarrow (C'; C'', \sigma', \omega', \gamma')} \quad \frac{}{(\mathbf{skip}; C, \sigma, \omega, \gamma) \rightsquigarrow (C, \sigma, \omega, \gamma)} \\
\\
\frac{(C, \sigma, \omega, \gamma) \rightsquigarrow \mathbf{abort}}{(C; C', \sigma, \omega, \gamma) \rightsquigarrow \mathbf{abort}} \\
\\
\frac{[B]_{\sigma_s} = \mathbf{tt}}{(\mathbf{if} (B) C_1 \mathbf{else} C_2, \sigma, \omega, \gamma) \rightsquigarrow (C_1, \sigma, \omega, \gamma)} \quad \frac{[B]_{\sigma_s} = \mathbf{ff}}{(\mathbf{if} (B) C_1 \mathbf{else} C_2, \sigma, \omega, \gamma) \rightsquigarrow (C_2, \sigma, \omega, \gamma)} \\
\\
\frac{[B]_{\sigma_s} = \mathbf{tt}}{(\mathbf{while} (B) C, \sigma, \omega, \gamma) \rightsquigarrow (C; \mathbf{while} (B) C, \sigma, \omega, \gamma)} \quad \frac{[B]_{\sigma_s} = \mathbf{ff}}{(\mathbf{while} (B) C, \sigma, \omega, \gamma) \rightsquigarrow (\mathbf{skip}, \sigma, \omega, \gamma)} \\
\\
\frac{l \notin \text{dom}(\sigma_h)}{(x := \mathbf{alloc}, \sigma, \omega, \gamma) \xrightarrow{\mathbf{alloc}(l)} (\mathbf{skip}, (\sigma_s[x \mapsto l], \sigma_h \uplus [l \mapsto 0]), \omega, \gamma)}
\end{array}$$

Figure 13: Annotated thread-local operational semantics.

Besides capturing a semantic relation between local states, primitive commands also have a syntactic representation (e.g. , the name of the command itself). We overload these notions in the rules to allow us flexibility in characterizing effects without having to choose a fixed command set *a priori*. It is straightforward to define an interpretation function that maps the nominal representation of a command with its relational definition.

The effect of evaluating primitive command  $c$  in state  $\langle c, \sigma, \omega, \gamma \rangle$  is a new state  $\langle \mathbf{skip}, \sigma', \omega, \gamma \rangle$  if  $(\sigma, \sigma') \in c$ . If command  $c$  can effect no transition from local state  $\sigma$ , the thread aborts.

A thread may fork a child. We associate a syntactic assertion  $P$  with a **fork** command that captures the resources required by the child thread. (The existence of a syntactic separation logic proof means we can always annotate a *fork* command with such assertions.) If the assertion is not satisfiable within the current local state (i.e. ,  $\sigma, \omega \not\models P * \mathbf{true}$ ), the thread aborts. Otherwise, we can partition the state, associating the resources (memory and channels) needed by the forked thread (denoted as  $\sigma'$  and  $\omega'$  in the rule) from those required by the parent. The effect annotation **fork** ( $t, C, \sigma', \omega'$ ) records this action. We assume that the fresh thread identifier  $t$  created on each **fork** invocation is strictly greater (with respect to  $<$ ) than previous ones.

Creating a new channel augments the set of channel identifiers for both **wait** and **grant** actions.

In order for a thread to **grant** resources to another thread, the assertion that defines the structure of these resources must hold in the current local state; note that the only resources that can be propagated along channels are locations and values in the thread's local state—channel identifiers can only be distributed at fork-time. If the assertion indeed holds, then the resources it describes can be transferred on the channel. This semantics ensures all concurrently executing threads operate over disjoint portions of shared memory. The semantics of **wait** transfers a set of resources to the executing thread whose composition is determined by the assertion on the corresponding **grant** that communicates on channel  $s$ . The rules for sequencing and loops are standard.

Memory allocation annotates the transition with an allocation label **alloc**( $l$ ) for fresh location  $l$ .

**Global semantics.** The global semantics is defined as a transition relation  $\Longrightarrow \in \mathcal{P}(\text{GState} \times \text{GState})$ , with **GState** defined as follows:

$$\begin{aligned} \delta &\in \text{TMap} : \text{Tid} \rightarrow \text{TState} \\ \eta &\in \text{CMap} : \text{Cid} \rightarrow \text{Heap} \uplus \{\Delta, \nabla\} \\ \kappa &\in \text{GState} : \text{TMap} \times \text{CMap} \times \mathcal{P}(\text{Loc}) \end{aligned}$$

The rules of the global semantics are given in Fig. 14. Every transition in the global semantics corresponds to a transition for some thread  $t$  in the local semantics. The global semantics models the pool of active threads with a thread-map in **TMap**, and models the resources held by channels with a channel-map in **CMap**. The set of locations in  $\mathcal{P}(\text{Loc})$  represents unallocated locations in the heap. Because we never return locations to this set, all allocated locations are fresh.

Given a channel map  $\eta \in \text{CMap}$ , a given channel  $s \in \text{dom}(\eta)$  can be either allocated but unused, denoted by  $\eta(s) = \Delta$ , in use, denoted by  $\eta(s) \in \text{Heap}$ , or finalised, denoted by

$$\begin{array}{c}
\frac{(C, \sigma, \omega, \gamma) \rightsquigarrow (C', \sigma', \omega', \gamma')}{([t \mapsto C, \sigma, \omega, \gamma] \uplus \delta, \eta, \mathcal{L}) \Longrightarrow ([t \mapsto C', \sigma', \omega', \gamma'] \uplus \delta', \eta, \mathcal{L})} \\
\\
\frac{(C, \sigma, \omega, \gamma) \overset{\text{fork}(t_2, C_2, \sigma_2, \omega_2)}{\rightsquigarrow} (C', \sigma', \omega', \gamma')}{([t_1 \mapsto C, \sigma, \omega, \gamma] \uplus \delta, \eta, \mathcal{L}) \Longrightarrow ([t_1 \mapsto C', \sigma', \omega', \gamma'] \uplus [t_2 \mapsto C_2, \sigma_2, \omega_2, \emptyset] \uplus \delta, \eta, \mathcal{L})} \\
\\
\frac{(C, \sigma, \omega, \gamma) \overset{\text{newch}(s)}{\rightsquigarrow} (C', \sigma', \omega', \gamma') \quad s \notin \text{dom}(\eta)}{([t \mapsto C, \sigma, \omega, \gamma] \uplus \delta, \eta, \mathcal{L}) \Longrightarrow ([t \mapsto C', \sigma', \omega', \gamma'] \uplus \delta, \eta \uplus [s \mapsto \Delta], \mathcal{L})} \\
\\
\frac{(C, \sigma, \omega, \gamma) \overset{\text{grant}(s, \sigma)}{\rightsquigarrow} (C', \sigma', \omega', \gamma') \quad \eta(s) = \Delta}{([t \mapsto C, \sigma, \omega, \gamma] \uplus \delta, \eta, \mathcal{L}) \Longrightarrow ([t \mapsto C', \sigma', \omega', \gamma'] \uplus \delta, \eta[s \mapsto \sigma], \mathcal{L})} \\
\\
\frac{(C, \sigma, \omega, \gamma) \overset{\text{wait}(s, \sigma)}{\rightsquigarrow} (C', \sigma', \omega', \gamma') \quad \eta(s) = \sigma}{([t \mapsto C, \sigma, \omega, \gamma] \uplus \delta, \eta, \mathcal{L}) \Longrightarrow ([t \mapsto C', \sigma', \omega', \gamma'] \uplus \delta, \eta[s \mapsto \nabla], \mathcal{L})} \\
\\
\frac{(C, \sigma, \omega, \gamma) \overset{\text{alloc}(l)}{\rightsquigarrow} (C', \sigma', \omega', \gamma') \quad l \in \mathcal{L}}{([t \mapsto C, \sigma, \omega, \gamma] \uplus \delta, \eta, \mathcal{L}) \Longrightarrow ([t \mapsto C', \sigma', \omega', \gamma'] \uplus \delta, \eta, \mathcal{L} \setminus \{l\})} \\
\\
\frac{(C, \sigma, \omega, \gamma) \rightsquigarrow \mathbf{abort}}{([t \mapsto C, \sigma, \omega, \gamma] \uplus \delta, \eta, \mathcal{L}) \Longrightarrow \mathbf{abort}}
\end{array}$$

Figure 14: Annotated global operational semantics.

$\eta(s) = \nabla$ . Once finalised, a channel identifier cannot be reused (fortunately we have an infinite supply of fresh channel identifiers). The semantics enforces consistency between calls to **wait**, **grant** and **newchan**. A thread can only take a local transition acquiring or releasing a resource if it is consistent with the global channel map.

Thread-local steps are mirrored in the global semantics by updating the state of the thread in the thread map. A fork operation augments the thread map by associating the thread identifier with a new thread-local configuration whose components are specified on the local transition. Creating a new channel is permissible only if the channel identifier does not already exist. A **grant** action is allowed if the channel is unused; after the action the channel is associated with the heap resources provided by the **grant**. A **wait** action is allowed if the resources demanded by the **wait** are provided by the channel (via a **grant**); if so, the channel becomes finalised. Allocation removes the location from the set of available locations.

**Definition 8.** For a global state  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$  and thread  $t \in \text{dom}(\delta)$  such that  $\delta(t) = \langle C, \sigma, (\omega_w, \omega_g), \gamma \rangle$ , we define  $\text{waits}(\kappa, t) \triangleq \omega_w$  and  $\text{grants}(\kappa, t) \triangleq \omega_g$ .

**Definition 9** (well-formedness). A global state  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$  is well-formed if:

1. The set of all heap locations either allocated and referred to in  $\delta$  and  $\eta$  is disjoint from  $\mathcal{L}$ , the set of unallocated heap locations.
2. The thread-local states in  $\delta$  and states stored in the channel map  $\eta$  can be joined using  $\oplus$  to give a well-defined state. Note that, as  $\oplus$  is associative and commutative,

this state is unique, and also that any pair of states from  $\delta$  and  $\eta$  can be joined to give a well-defined result.

3. For any channel  $c \in \text{Cid}$ , there exists at most one thread  $t_1$  such that  $c \in \text{grants}(\kappa, t_1)$  and at most one thread  $t_2$  such that  $c \in \text{waits}(\kappa, t_2)$ . If  $c \in \text{grants}(\kappa, t_1)$ , then  $\eta(c) = \Delta$ . If  $c \in \text{waits}(\kappa, t_2)$ , then  $\eta(c) \in \text{Heap} \uplus \{\Delta\}$ .

Well-formedness expresses fundamental linearity assumptions on states and channels; essentially, it ensures elements of the state can be owned by only one thread at once.

**Lemma 10.** *The global transition relation  $\Longrightarrow$  preserves well-formedness.*

**Assumption 1.** *We assume all global states are well-formed, unless explicitly stated otherwise.*

**Definition 11** (trace). *A trace is a (finite or infinite) sequence of global states  $\mathcal{K} = \kappa_0 \kappa_1 \dots$  such that for every  $i$ ,  $\kappa_i \Longrightarrow \kappa_{i+1}$ . A finite (infinite) trace is called terminating (nonterminating).*

**Definition 12** (child thread, thread order). *Each local state  $\langle C, \sigma, \omega, \gamma \rangle$  includes the set  $\gamma$  of forked child thread identifiers. For convenience, we define  $\text{child}(\kappa, t)$  as the set of children for thread  $t$  in global state  $\kappa$ .*

*The children of a thread are ordered (under relation  $<$ ), with the ordering following the order in which they were created. That is, suppose we have a trace  $\mathcal{K}$  with initial state  $\langle \delta, \eta, \mathcal{L} \rangle$  such that  $c_1, c_2 \notin \text{dom}(\delta)$ . Given an arbitrary state  $\kappa_i \in \mathcal{K}$ , if  $\{c_1, c_2\} \in \text{child}(\kappa, t)$  and  $c_1 < c_2$ , then  $c_1$  was created earlier in the trace than  $c_2$ .*

We denote by  $\Longrightarrow^*$  the reflexive transitive closure of the global transition relation. We sometimes write ' $\xrightarrow{t}$ ' to denote a global transition resulting from the thread  $t$  taking a step, and call it a transition over thread  $t$ .

**Definition 13** (sequentialised trace). *We define the sequentialised transition relation  $\Longrightarrow_s \subseteq \Longrightarrow$  as the transition relation in which a step  $\kappa \xrightarrow{t} \kappa'$  can only be taken if all threads in  $\text{child}(\kappa, t)$  have reduced to **skip**. We say that a trace  $\mathcal{K}$  is sequentialised with respect to  $t$  if every transition over thread  $t$  is in  $\Longrightarrow_s$ . We denote this transition relation by  $\Longrightarrow_{s(t)}$ .*

## B Soundness of Analysis

In this section we prove the soundness of our analysis, i.e., that our parallelising transformation does not introduce any new behaviours to the system. The proof of the behaviour preservation (Theorem 25) is based on a simulation argument: each trace of the parallelised program must correspond to a trace of the sequential program and preserve as an invariant a given binary relation between states of the parallelised and the sequential program. The proof is structured as follows:

1. We establish a simulation invariant between the parallelised program and the sequential program in two steps: first for a program decorated with calls to **grant**, **wait** and **newch** and a program with these calls erased (Lemma 14), and then

for a program additionally decorated with **fork** calls and a **fork**-erased program (Lemma 15). The combined invariant relates every nonfaulting sequentialised trace of the fully decorated program with a trace of the sequential program. By *sequentialised*, we mean that forked children must execute to completion before their parent threads can be scheduled (Def. 13).

2. We show that, under the assumption that forked child threads never wait for channels granted by their parent or later-forked child threads, any terminating nonfaulting trace of the parallelised program can be reordered into a sequentialised trace with the same thread-local behaviour. This is established by showing that traces of the program parallelised by our analysis have signalling barriers aligned with the thread ordering (Lemma 21), and that any such trace can be reordered into a sequentialised trace (Lemma 20).
3. Previous steps establish good behaviour for *non-aborting* traces. Inserting **grant**, **wait**, **newch** and **fork** may introduce aborts. However, we also verify our mutated program in separation logic. This establishes, for every state satisfying the program precondition, that the program cannot abort.

We start with a lemma that establishes a simulation invariant between a program decorated with calls to **grant**, **wait** and **newch** and a program with these calls replaced with **skip**.

**Lemma 14** (**wait**, **grant**, **newch** insertion). *Let  $\llbracket - \rrbracket_{\text{ch}}^{\downarrow}$  be a program transformation which replaces calls to **newch**, **grant** and **wait** with **skip**. Let  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$  and  $\kappa'$  be global states. We say that the invariant  $I(t, \kappa, \kappa')$  is satisfied if:*

- The thread identifier  $t$  is in  $\text{dom}(\delta)$ .
- For  $\delta(t) = \langle C, \sigma, \omega, \gamma \rangle$ , the program  $C$  executed by  $t$  does not contain **fork**.
- The thread waits for all channels it grants on, i.e.,  $\omega_g \subseteq \omega_w$ , and all channels in  $\omega_w \setminus \omega_g$  have already been granted on but are not yet finalised, i.e.,  $\forall c \in \omega_w \setminus \omega_g. \eta(c) \notin \{\nabla, \Delta\}$ .
- Let  $\sigma_{\text{ch}} \triangleq \bigotimes_{c \in \omega_w \setminus \omega_g} \eta(c)$  be all resources the thread does not grant but waits for. Then there exists  $\eta' \subseteq \eta$  such that  $\kappa' = \langle \delta[t \mapsto \langle \llbracket C \rrbracket_{\text{ch}}^{\downarrow}, \sigma \oplus \sigma_{\text{ch}}, (\emptyset, \emptyset), \gamma \rangle], \eta', \mathcal{L} \rangle$  and for all channels  $c$  held by threads other than  $t$ ,  $\eta'(c) = \eta(c)$ .

For any states  $\kappa, \kappa', \kappa_2$ , if  $I(t, \kappa, \kappa')$  and  $\kappa \Vdash \kappa_2$ , then there exists a state  $\kappa'_2$  such that  $\kappa' \Vdash^* \kappa'_2$  and  $I(t, \kappa_2, \kappa'_2)$ . This property can be alternatively represented by the following diagram:

$$\begin{array}{ccc}
 \kappa & \xrightarrow{I(t)} & \kappa' \\
 \Downarrow & & \Downarrow^* \\
 \kappa_2 & \xrightarrow{I(t)} & \kappa'_2
 \end{array}$$

Intuitively, the invariant  $I$  expresses the simulation relation between programs containing **wait**, **grant** and **newch** (left-hand side of the diagram), and the equivalent program without them (right-hand side).

*Proof.* There are two cases: either  $\kappa \Longrightarrow \kappa_2$  is a transition over  $t$ , or it is a transition over some other thread  $t' \neq t$ . If  $\kappa \Longrightarrow \kappa_2$  is a transition over  $t$ , we proceed by structural induction over  $C$ .

- $C$  is a primitive command in **Prim**. In this case, the result holds by the assumption of behavioural monotonicity for primitive commands.
- $C$  is a **grant** or **wait** for some channel  $c$ . In both of these cases, the corresponding command  $\llbracket C \rrbracket_{\text{ch}}^\downarrow$  will be **skip**. Consequently the transition  $\kappa' \Longrightarrow \kappa'_2$  will be a  $\tau$ -transition, meaning  $\kappa' = \kappa'_2$ .

By assumption, the join of the thread-local state  $\sigma$  and state  $\sigma_{\text{ch}}$  stored in accessible channels in  $\kappa$  is equal to the thread-local state in  $\kappa'$ . Calling a **grant** takes some local state and pushes it into a channel, which preserves this property. Similarly calling a **wait** pulls some state out of an accessible channel and into local state, which also preserves the property.

- $C$  is a **newch**. In this case, the corresponding command also be a **skip**, so  $\kappa' = \kappa'_2$ . The effect of **newch** will be to initialise a fresh channel not already in  $\eta$ , and add these channels to the channel state for  $\kappa_2$ . These channels are erased by the invariant, so the property is preserved.
- $C$  is **alloc**. The set  $\mathcal{L}$  of unallocated locations is the same in  $\kappa$  and  $\kappa'$ . Consequently, if a location can be allocated in  $\kappa$ , it can also be allocated in  $\kappa'$ .
- $C$  is a sequential composition, loop, or conditional. The result follows trivially by appeal to the induction hypothesis.

Suppose now that the transition is over some other thread  $t' \neq t$ . By assumption (the third property of  $I$ ), we know that the channels erased by the invariant are not shared with any other thread. We can therefore establish straightforwardly that such a transition must be replicated exactly between  $\kappa'$  and  $\kappa'_2$ . The invariant is not disturbed, because it does not mutate the content of other threads.  $\square$

The following lemma establishes a simulation invariant between a program containing calls to **fork** and a program with these calls replaced with **skip**.

**Lemma 15 (fork insertion).** *Let  $\llbracket - \rrbracket_{\text{fk}}^\downarrow$  be a program transformation which replaces calls to **fork** with **skip**. We say that the invariant  $J(t, \kappa, \kappa')$  is satisfied if:*

- The thread identifier  $t$  is in  $\text{dom}(\delta)$ .
- For  $\delta(t) = \langle C, \sigma, \omega, \gamma \rangle$ , any calls to **fork** in  $C$  do not themselves include calls to **fork**.
- Let  $\sigma_{\text{fk}} = \bigotimes \{ \sigma_c \mid c \in \text{child}(t) \wedge \delta(c) = \langle C_c, \sigma_c, \omega_c, \gamma_c \rangle \}$  be thread-local states and  $\omega_{\text{fk}} = \bigoplus \{ \omega_c \mid c \in \text{child}(t) \wedge \delta(c) = \langle C_c, \sigma_c, \omega_c, \gamma_c \rangle \}$  channel states of forked child threads conjoined. By the well-formedness condition on global states (Assumption 9), both of these must be well-defined.
- There exists at most one child thread  $c \in \text{child}(t)$  such that  $\delta(c) = \langle C_c, \sigma_c, \omega_c, \gamma_c \rangle$  and  $C_c \neq \text{skip}$ .

- Either of the following two cases holds:

– If  $\forall c \in \text{child}(t). \delta(c) = \langle \mathbf{skip}, \sigma_c, \omega_c \rangle$ , then  $\kappa' = \langle \delta[t \mapsto \langle \llbracket C \rrbracket_{\text{fk}}^\downarrow, \sigma \oplus \sigma_{\text{fk}}, \omega \uplus \omega_{\text{fk}} \rangle], \eta, \mathcal{L} \rangle$ .

(In other words, all child threads of  $t$  have terminated, and the main thread corresponds to the original sequential state.)

– If  $\exists! c \in \text{child}(t). \delta(c) = \langle C_c, \sigma_c, \omega_c, \gamma_c \rangle \wedge C_c \neq \mathbf{skip}$ , then  $C = \mathbf{skip}; C_1$  and  $\kappa' = \langle \delta[t \mapsto \langle C_c; \llbracket C_1 \rrbracket_{\text{fk}}^\downarrow, \sigma \oplus \sigma_{\text{fk}}, \omega \uplus \omega_{\text{fk}}, \emptyset \rangle], \eta, \mathcal{L} \rangle$ .

(In other words, there exists exactly one unterminated child thread, and join of the states of the child threads and main thread correspond to the original sequential state.)

For any states  $\kappa, \kappa', \kappa_2$ , if  $J(t, \kappa, \kappa')$ , and  $\kappa \Longrightarrow_{s(t)} \kappa_2$ , then there exists a state  $\kappa'_2$  such that  $\kappa' \Longrightarrow^* \kappa'_2$  and  $J(t, \kappa_2, \kappa'_2)$ . This property can be alternatively represented by the following diagram:

$$\begin{array}{ccc}
 \kappa & \xrightarrow{J(t)} & \kappa' \\
 \Downarrow & & \Downarrow \\
 \textcircled{s} & & * \\
 \kappa_2 & \xrightarrow{J(t)} & \kappa'_2
 \end{array}$$

The invariant  $J$  expresses the simulation relation between programs containing **fork** (left-hand side of the diagram) and the equivalent program without it (right-hand side).

*Proof.* Unlike with Lemma 14, there are *three* cases: either  $\kappa \Longrightarrow_{s(t)} \kappa_2$  is a transition over  $t$ ; or it is a transition over some member of the set  $\text{child}(t)$ ; or the transition is over some other, unrelated thread. First suppose that the transition is over  $t$ . We proceed by structural induction on  $C$ :

- $C \in \text{Prim}$ . By assumption, any behaviour in a small state  $\sigma$  will be reflected in a big one, which suffices to ensure that the invariant holds.
- $C$  is **alloc**. The set  $\mathcal{L}$  of unallocated locations is the same for  $\kappa$  and  $\kappa'$ . By the same argument as Lemma 14, the same allocation step is therefore possible in both global states.
- $C$  is **wait** or **grant**. By the well-formedness assumption, resources held by channels must be disjoint from those held by other threads. This suffices to ensure that **wait** and **grant** transitions can take place identically in the corresponding **fork**-free program.
- $C$  is **fork**. This will correspond to a  $\tau$ -transition in the **fork**-erased program. Because the transition over the program with **fork** is sequentialised with respect to  $t$ , the transition can only be over  $t$  if all the children of  $t$  have reduced to **skip**. Calling **fork** pushes some state to the child thread, and results in exactly one active (non-**skip**) child thread.
- $C$  is a loop, conditional, sequential composition. Property holds by appeal to the induction hypothesis.

Now suppose the transition is over some thread  $c \in \mathbf{child}(t)$ . By assumption (the fourth property of  $J$ ), there can be at most one such thread. The state held by the child thread must be a substate of that held in the **fork**-free program. Consequently, an almost identical structural induction argument suffices to show that the simulation property holds.

Finally, suppose that the transition is over some other thread  $t'$  that is neither  $t$  nor a child of  $t$ . Such threads cannot be affected by the behaviour of  $t$  and its children. It is straightforward to show that any transition for the program can take place in the corresponding **fork**-free program.  $\square$

The previous two lemmas combined establish an invariant relating nonfaulting sequentialised traces of the parallelised program with traces of the sequential program. It remains to show that terminating nonfaulting traces of the parallelised program can be reordered into sequentialised traces with the same thread-local behaviour. We do this by introducing a well-founded order on traces in which sequentialised traces are minimal elements, and showing that each signal-ordered trace (i.e., having its signalling barriers aligned with the thread ordering) can be reordered to a behaviourally equivalent trace smaller with respect to the well-founded order, eventually leading to a behaviourally equivalent sequentialised trace.

**Definition 16.** *We say that traces  $\mathcal{K}_1$  and  $\mathcal{K}_2$  are behaviourally equivalent if, for each thread identifier  $t$ , the sequence of transitions over  $t$  in  $\mathcal{K}_1$  and  $\mathcal{K}_2$  are identical.*

We first show that consecutive transitions of two different threads can be reordered if the second thread does not wait for any channel that the first thread grants on. We use this property in the subsequent lemma to inductively construct a behaviourally equivalent sequentialised trace.

**Lemma 17** (trace reordering). *Suppose we have a two-step trace:*

$$\mathcal{K} = \langle \delta_1, \eta_1, \mathcal{L}_1 \rangle \xrightarrow{t_1} \langle \delta_2, \eta_2, \mathcal{L}_2 \rangle \xrightarrow{t_2} \langle \delta_3, \eta_3, \mathcal{L}_3 \rangle$$

*such that  $t_1 \neq t_2$ , and that  $t_1, t_2 \in \text{dom}(\delta_1)$ . Let  $\delta_1(t_1) = \langle C, \sigma, \omega, \gamma \rangle$  and  $\delta_1(t_2) = \langle C', \sigma', \omega', \gamma' \rangle$ . Suppose that  $t_2$  does not wait for any channel that  $t_1$  grants on, i.e., that  $\omega_g \cap \omega'_w = \emptyset$  holds. Then there exists a thread environment  $\eta'$ , an unallocated set  $\mathcal{L}'$ , and a two-step trace:*

$$\mathcal{K}' = \langle \delta_1, \eta_1, \mathcal{L}_1 \rangle \xrightarrow{t_2} \langle \delta_1[t_2 \mapsto \delta_3(t_2)], \eta', \mathcal{L}' \rangle \xrightarrow{t_1} \langle \delta_3, \eta_3, \mathcal{L}_3 \rangle$$

*The new trace  $\mathcal{K}'$  is behaviourally equivalent (Def. 16) to  $\mathcal{K}$ .*

*Proof.* We proceed by case-splitting on the shape of transitions over  $t_1$  and  $t_2$ :

- One or both of the transition are thread-local (i.e. do not result in a labelled transition in the thread-local semantics). It is straightforward to see that neither transition can affect the other one, and the rearrangement result follows trivially.
- One or both of the transitions is a **fork**. The effect of such a transition is localised to the parent and newly-created child thread. By assumption, neither thread was created by either of these transitions. The rearrangement result follows straightforwardly.



- One or both transitions is a **newch**. Channels can only be transferred by **fork**, and by assumption, neither transition created either of the threads. Consequently the effect of a **newch** call is confined to the local thread.
- Both transitions are **wait** (resp. **grant**). The threads could only affect one another if both waited (resp. granted) on the same channel, but this is ruled out by the well-formedness assumption that each end of each channel is held by at most one thread.
- One transition is a **wait** and the other is a **grant**. Once again, the threads could only affect one another if they waited and granted on the same channel. However, by assumption of the lemma, the thread  $t_1$  cannot grant on any channel on which  $t_2$  can wait. By the structure of the semantics, the transition  $t_1$  cannot wait for a channel that is subsequently granted. So this case cannot occur.
- One or both transitions are a **alloc**. If one operation is a **alloc** and the other another command, the threads trivially cannot affect one another. If both transitions are **alloc**, the locations allocated must be distinct, and the transitions again can be exchanged straightforwardly.

The constructed trace  $\mathcal{K}'$  is trivially behaviourally equivalent to  $\mathcal{K}$ , since the thread-local actions for each thread are identical.  $\square$

**Definition 18** (signal order). *We describe a trace  $\mathcal{K}$  as signal-ordered with respect to  $t$  if  $t$  is a thread identifier such that for all states  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle \in \mathcal{K}$ :*

$$\forall c \in \text{child}(\kappa, t). \forall k \in \text{waits}(\kappa, c). \\ ((\exists c' \in \text{child}(\kappa, t). c' < c \wedge k \in \text{grants}(\kappa, c')) \vee (\eta(k) \notin \{\perp, \nabla, \Delta\}))$$

*(In other words, if a child of  $t$  can wait on a channel  $k$ , some earlier child must be able to grant on it, or the channel must already be filled.)*

**Definition 19** (degree of sequentialisation). *Let  $\mathcal{K}$  be a trace. If  $\mathcal{K}$  is not sequentialised with respect to  $t$ , let  $c \in \text{child}(t)$  be the earliest thread in the child order with an unsequentialised transition (that is, for all other children  $c'$  with unsequentialised transitions,  $c' > c$ ). Now, let  $F$  be the number of transitions from the originating **fork** transition for  $c$  and the end of the trace. Let  $K$  be the number of transitions over  $c$  after the first unsequentialised transition. Let  $D$  be the number of intervening transitions between the originating **fork** and the first unsequentialised transition.*

*We call such a tuple  $(F, K, D)$  the degree of sequentialisation with respect to  $t$ . We order such tuples lexicographically, with left-to-right priority, and lift this order to give a partial order on traces  $<_t$ . This order on traces is well-founded, with sequentialised traces occupying the minimum position in the order.*

Using the ordering on traces introduced in Definition 19 we now show that a trace that is signal-ordered with respect to a thread  $t$  can be reordered to a behaviourally equivalent trace that is sequentialised with respect to  $t$ .

**Lemma 20** (sequentialisation). *Let  $\mathcal{K}$  be a trace such that*

- $\mathcal{K}$  is signal-ordered with respect to  $t$ ; and
- Let  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$  be the final state in  $\mathcal{K}$  and let  $t'$  be the maximum thread in  $\text{child}(\kappa, t)$  with transitions in  $\mathcal{K}$ . For all threads  $t'' \in \text{child}(\kappa, t)$  such that  $t'' < t'$ ,  $\delta(t'') = \langle \mathbf{skip}, \sigma, \omega, \gamma \rangle$  for some  $\sigma, \omega, \gamma$ . (Note this holds automatically if  $\mathcal{K}$  is terminating.)

Then there must exist a trace  $\mathcal{K}'$  sequentialised with respect to  $t$  such that  $\mathcal{K}$  is behaviourally equivalent to  $\mathcal{K}'$ .

*Proof.* We show that any out-of-order trace  $\mathcal{K}$  can be reordered by applying Lemma 17, to give a behaviourally equivalent trace with a lower degree of sequentialisation w.r.t.  $t$ .

Consider a trace  $\mathcal{K}$  that is signal-ordered but not sequentialised. Let  $t_2$  be the earliest thread in the child order for  $t$  that has an unsequentialised transition. By the assumption that the transition is unsequentialised, the preceding transition in the trace must be on some other thread  $t_1 \neq t_2$ , and must not be the **fork** that created the out-of-order transition. Note that  $t_1$  also must not be some child of  $t$  earlier in the child order—otherwise the transition over  $t_1$  would be the earliest out-of-order transition.

By the definition of signal-ordering we know that  $\text{wait}(\kappa, t_2) \cap \text{grant}(\kappa, t_1) = \emptyset$ . Consequently, we can apply Lemma 17, and push the transition over  $t_2$  earlier than the transition over  $t_1$ . Call the resulting trace  $\mathcal{K}''$ .

By Lemma 17,  $\mathcal{K}''$  is behaviourally equivalent to  $\mathcal{K}$ . As a trivial consequence,  $\mathcal{K}''$  is also signal-ordered. As a deeper consequence,  $\mathcal{K}'' <_t \mathcal{K}$  (with respect to the order given in Def. 19). There are three cases:

- The transition shifts left, but is still unsequentialised. The number of intervening transitions from the originating **fork** decreases, while the other two measures are unchanged.
- The left shift means the transition is sequentialised, but more transitions over the same thread are unsequentialised. The number of transitions from the first unsequentialised transition decreases, while the distance from the initial **fork** to the end of the thread is unchanged.
- The left shift means all transitions for the thread are sequentialised. The next target thread must be later in the child order, meaning it is later in the trace. The number of transitions from the initial **fork** to the end of the thread decreases.

As the order is well-founded, by repeatedly applying the lemma, we get a behaviourally-equivalent trace that is sequentialised w.r.t.  $t$ .  $\square$

We now establish that the pattern of barriers that we insert into the program during our parallelisation analysis results in a trace that is signal-ordered. Recall from §5 that the parallelised program constructed by our analysis has the following form<sup>7</sup>:

---

<sup>7</sup>Recall also that we define **for** in terms of **while**.

```

ci := newch();
grant(ci);
C1;
while(B) {
  C2;
  cj := newch();
  fork[P](C3);
  ci := cj;
}
wait(ci);

```

Here we assume that we only need a single active pair of channels  $ci/cj$ —the argument generalises straightforwardly to the case with  $n$  channels. In addition, we assume the sub-programs  $C1$  and  $C2$  do not include `wait`, `grant`, `newch` and `fork`, and that they do not assign to the variables  $ci$  and  $cj$ .

The assertion  $P$  is defined so that:  $P \vdash \text{fut}(ci, P_1) * \text{req}(cj, P_2) * F$ , for some  $P_1, P_2$  and  $F$ ; and  $F \not\vdash \text{fut}(x, P')$  and  $F \not\vdash \text{req}(x, P')$  for any  $x$  and  $P'$ . (In other words, the forked thread can only wait for  $ci$  and grant on  $cj$ .)

**Lemma 21** (construction of signal-ordered traces). *Let  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$  be a state and  $t$  a thread identifier such that  $\delta(t) = \langle C, \sigma, (\emptyset, \emptyset), \emptyset \rangle$  and  $C$  matches the standard form for parallelised programs, described above. Let  $\mathcal{K}$  be a trace with initial state  $\kappa$ . Then  $\mathcal{K}$  is signal-ordered with respect to  $t$ .*

*Proof.* We establish this property by defining an invariant  $L(\kappa)$ :

- Let  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$  and  $\delta(t) = \langle C, \sigma, \omega, \gamma \rangle$ .
- There exists a channel  $c_1$  such that  $\sigma_s(ci) = c_1$ , and  $\text{waits}(\kappa, t) = \{c_1\}$ .
- Either  $\eta(c_1) \notin \{\perp, \nabla, \Delta\}$ , or  $\exists t' \in \text{child}(\kappa, t). c_1 \in \text{grants}(\kappa, t')$ .

We now show that signal-ordering holds for every step of the semantics, and that the invariant holds for any state immediately preceding the execution of an iteration of the main `while`-loop; that is, at any point when  $C = \text{while}(B)\{\dots\}; \text{wait}(ci);$ .

- Suppose we execute from the beginning of the program (the first command in  $C1$ ) to the beginning of the first iteration `while`. It is simple to see by inspection of the semantics that the resulting state  $\kappa$  will satisfy the invariant  $L$ . We create the channel  $c_1$ , and the call to `grant` ensures  $\eta(c_1)$  is associated with some state, as required by the invariant.
- Now suppose we execute a single iteration of the loop: that is, we assume the boolean condition  $B$  holds, and apply the appropriate rule in the operational semantics. The invariant cannot be disturbed by other running threads, because by definition they can only call `grant` on the channel  $c_1$ .

By assumption, the sub-program  $C2$  cannot disturb the invariant. We call `newch`, which creates a channel  $c_2$ , associated with the variable  $cj$ . Now we observe that, by the structure of  $P$  call to `fork` creates a thread  $t'$  with  $\text{waits}(\kappa, t') = \{c_1\}$  and  $\text{grants}(\kappa, t') = \{c_2\}$ . As a result, in the resulting state  $\text{waits}(\kappa, t) = \{c_2\}$ . The final assignment associates  $c_2$  to  $ci$ , which reestablishes the invariant.

Each step in this execution satisfies signal-ordering. By induction this completes the proof.  $\square$

In Lemma 21 we assumed a specific program form constructed by the parallelising transformation from §5.2. However, the soundness proof could be adapted to hold for other parallelisation backends. As long as the analysis generates a parallelised program with traces that are signal-ordered we could establish an analogue of Lemma 21 and proceed with other steps of the soundness proof in the same way.

**Lemma 22** (synchronisation erasure). *Let  $C$  be a sequential program, and  $C_{\text{par}}$  be a program resulting from applying our parallelisation analysis. Then  $C = \llbracket \llbracket C_{\text{par}} \rrbracket_{\text{fk}}^{\downarrow} \rrbracket_{\text{ch}}^{\downarrow}$ , where the (syntactic) equality of programs is defined up to insertion/erasure of **skip** statements — that is syntactic manipulation using the identity  $C = C; \text{skip}$ .*

*Proof.* Our analysis only inserts channels and calls to fork. The result follows trivially.  $\square$

**Lemma 23** (materialised variables). *Materialising local variables, as described in §3, has no effect on the operational behaviour of a program.*

*Proof.* Let  $C$  be a program, and let  $C'$  be a corresponding program in which the local variables are materialised. Let  $\mathcal{K}$  be a trace in which the program  $C$  is executed. When inserted, the materialised variables are written but not read (they are *subsequently* used in conditionals after loop-splitting). Therefore, there must be a trace  $\mathcal{K}'$  in which  $C'$  is substituted for  $C$ , which is identical to  $\mathcal{K}$  aside from the materialised variables.  $\square$

**Lemma 24** (loop-splitting). *Loop-splitting, as described in §5.2, has no effect on the operational behaviour of the program.*

*Proof.* By the fact that either a conditional or its negation must hold, and a straightforward appeal to the semantics of loops.  $\square$

Finally, we prove the main theorem stating the soundness of our parallelisation analysis.

**Theorem 25** (parallelisation soundness). *Let  $C$  be a sequential program, and let  $C_{\text{par}}$  be the program resulting from applying our parallelisation analysis to  $C$  (including variable materialisation and loop-splitting). Let  $\mathcal{K}$  be a nonfaulting terminating trace with initial state  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$ , and let  $t$  be a thread identifier such that  $\delta(t) = \langle C_{\text{par}}, \sigma, (\emptyset, \emptyset), \emptyset \rangle$ . Then there exists a trace  $\mathcal{K}''$  with initial state  $\kappa'' = \langle \delta[t \mapsto \langle C, \sigma, (\emptyset, \emptyset), \emptyset \rangle], \eta, \mathcal{L} \rangle$  such that:*

1. *For all thread identifiers  $t'$  defined in  $\mathcal{K}''$  such that  $t \neq t'$ , their thread-local behaviour in  $\mathcal{K}$  is identical to  $\mathcal{K}''$ .*
2. *Let  $\kappa_{\Omega}$  be the final state in  $\mathcal{K}$ , and  $\kappa_{\Omega}''$  the final state of  $\mathcal{K}''$ . Let  $\sigma$  be the local state associated with thread  $t$  in  $\kappa_{\Omega}$  and  $\sigma''$  the corresponding state associated with  $t$  in  $\kappa''$ . There exists a state  $\sigma'$  such that  $\sigma \oplus \sigma' = \sigma''$ .*

*Proof.* By Lemma 21 the trace  $\mathcal{K}$  must be signal-ordered with respect to  $t$ . Therefore by Lemma 20 there exists a behaviourally-equivalent trace  $\mathcal{K}_{\text{seq}}$  that is sequentialised with respect to  $t$ . Because  $t$  has no child threads in  $\kappa$ , we can choose  $\mathcal{K}_{\text{seq}}$  such that  $\kappa$  is also its initial state.

We now show that there exist traces  $\mathcal{K}'$  and  $\mathcal{K}''$ , such that transitions in  $\mathcal{K}$  are related to  $\mathcal{K}'$  by the invariant  $J(t)$  (Lemma 15) and transitions in  $\mathcal{K}'$  are related to  $\mathcal{K}''$  by the invariant  $I(t)$  (Lemma 14).

We proceed by induction on the length of a prefix of  $\mathcal{K}_{\text{seq}}$ . By Lemma 22,  $C = \llbracket \llbracket C_{\text{par}} \rrbracket_{\text{fk}}^{\downarrow} \rrbracket_{\text{ch}}^{\downarrow}$ . From this, it is straightforward to see that there exists a state  $\kappa' = \langle \delta[t \mapsto \langle \llbracket C_{\text{par}} \rrbracket_{\text{fk}}^{\downarrow}, \sigma, (\emptyset, \emptyset), \emptyset \rangle], \eta, \mathcal{L} \rangle$  such that  $J(t, \kappa, \kappa')$  and  $I(t, \kappa', \kappa'')$ .

Assume traces  $\mathcal{K}'$  and  $\mathcal{K}''$  exist for the first  $n$  transitions of  $\mathcal{K}_{\text{seq}}$ . Let  $\kappa_n$  be the final state in this prefix. By assumption there exist states  $\kappa'_n$  and  $\kappa''_n$  such that  $J(t, \kappa_n, \kappa'_n)$  and  $I(t, \kappa'_n, \kappa''_n)$ . Now show that for the next transition  $\kappa_n \xrightarrow{s(t)} \kappa_{n+1}$  there exist transitions  $\kappa'_n \xrightarrow{*} \kappa'_{n+1}$  and  $\kappa''_n \xrightarrow{*} \kappa''_{n+1}$  such that  $J(t, \kappa_{n+1}, \kappa'_{n+1})$  and  $I(t, \kappa'_{n+1}, \kappa''_{n+1})$ . This can be visualised by the following diagram:

$$\begin{array}{ccccc}
\kappa_n & \xrightarrow{J(t)} & \kappa'_n & \xrightarrow{I(t)} & \kappa''_n \\
\Downarrow \textcircled{s} & & \Downarrow * & & \Downarrow * \\
\kappa_{n+1} & \xrightarrow{J(t)} & \kappa'_{n+1} & \xrightarrow{I(t)} & \kappa''_{n+1}
\end{array}$$

This is an immediate consequence of Lemmas 14 and 15.

By induction, this suffices to establish the existence of the trace  $\mathcal{K}''$ . Corresponding thread-local behaviour for threads  $t' \neq t$ , and corresponding final state for thread  $t$  follow immediately from the definitions of  $I(t)$  and  $J(t)$ .

This result also holds if our analysis performs variable materialisation and loop-splitting—to show this, we need only appeal to Lemma 23 and Lemma 24.  $\square$

**Corollary 26.** *If the chosen post-condition in the proof of the program is precise, then the result of parallelisation is an equality, not a sub-state.*

*Proof.* Consequence of the fact that at most one sub-state of a given state can satisfy a precise assertion.  $\square$

## B.1 Termination

In this section we prove that our analysis cannot introduce nontermination to the parallelised program.

**Lemma 27** (ensuring grant). *Let  $C_{\text{par}}$  be a program resulting from applying our parallelisation analysis to some sequential program  $C$ . Let  $\mathcal{K}$  be a trace with initial state  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$ , such that  $\delta(t) = \langle C_{\text{par}}, \sigma, (\emptyset, \emptyset), \emptyset \rangle$ . For all states in  $\kappa' = \langle \delta, \eta, \mathcal{L} \rangle \in \mathcal{K}$  and all children  $c \in \text{child}(\kappa', t)$  if  $c$  has terminated in  $\kappa'$  (i.e. reduced to **skip**) then  $\text{grants}(\kappa', c) = \emptyset$ .*

*Proof.* Consequence of the fact that in our analysis, each **grant** is injected along every control-flow path. Consequently, a terminating thread must call **grant** for every channel to which it has access.  $\square$

**Lemma 28** (trace extension). *Let  $\mathcal{K}$  be a trace that is signal-ordered with respect to  $t$ , and let  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$  be the final state of  $\mathcal{K}$ . Let  $t'$  be the minimum child of  $t$  such that  $\delta(t') = \langle C, \sigma, \omega, \gamma, \rangle$  and  $C \neq \text{skip}$ . Then there exists a nonfaulting transition  $\kappa \xrightarrow{t'} \kappa'$ .*

*Proof.* By examination of the thread-local semantics, we can see that only calls to **wait** can block without faulting. For all other commands, either the thread can take a step, or it faults immediately. Consequently, the lemma reduces to asking whether every call to **wait** in the thread  $t'$  can take a step.

**wait** can only block if some prior thread has not called **grant** on some channel. By the definition of signal-ordered (Def. 18), the channels for which  $t'$  can wait must either contain state, or must be held by some child thread earlier in the order. By assumption, we know that all earlier threads have terminated, so by Lemma 27, we know that no thread holds a pending grant. Consequently, the channel must contain state, and **wait** can take a step.  $\square$

**Lemma 29** (trace erasure). *Let  $\mathcal{K}$  be a trace that is signal-ordered with respect to  $t$ , and let  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$  be the final state of  $\mathcal{K}$ . Let  $t'$  be the maximum child of  $t$  in  $\text{dom}(\delta)$ . Then there exists a trace  $\mathcal{K}'$  such that  $\mathcal{K}'$  has no transitions over  $t'$ , and all threads other than  $t'$  have identical behaviour in  $\mathcal{K}$  and  $\mathcal{K}'$ .*

*Proof.* By induction over the number of transitions over  $t$  in the trace. identify the final transition over  $t'$  in  $\mathcal{K}$ ,  $\kappa_n \xrightarrow{t'} \kappa_{n+1}$ . We erase this transition, and then re-execute the remainder of the transitions from the trace. By the definition of signal-order (Def. 18) no other child thread can call **wait** on any channel held by  $t$ . Consequently, we must be able to run the same sequence of thread-local actions after erasing the transition. By applying this process, we can erase all transitions over  $t'$ .  $\square$

**Theorem 30** (termination). *Let  $C$  be a sequential program, and let  $C_{\text{par}}$  be a program resulting from applying our parallelisation analysis to  $C$ . If  $C$  is guaranteed to terminate, then so is  $C_{\text{par}}$ .*

*Proof.* If  $C$  is terminating, then for any initial state, there must be a maximum number of steps that the thread  $t$  running  $C$  can take. To prove the theorem, we prove the contrapositive result: given that  $C_{\text{par}}$  does not terminate, we can construct a trace for  $C_{\text{seq}}$  where  $t$  takes more than this number of steps.

Suppose we have an initial state  $\kappa = \langle \delta, \eta, \mathcal{L} \rangle$  such that  $\delta(t) = \langle C_{\text{par}}, \sigma, (\emptyset, \emptyset), \emptyset \rangle$ . Suppose that  $\kappa$  can result in a nonterminating trace over  $t$ —that is, given any arbitrary bound  $n$ , there exists a trace with initial state  $\kappa$  and more than  $n$  steps over  $t$  and its child threads.

Construct a sequential state  $\kappa_{\text{seq}} = \langle \delta[t \mapsto \langle C_{\text{par}}, \sigma, (\emptyset, \emptyset), \emptyset \rangle], \eta, \mathcal{L} \rangle$ , and identify a bound  $q$  such that no trace starting with  $\kappa''$  can take more than  $q$  steps over  $t$ . Choose a trace  $\mathcal{K}$  with initial state  $\kappa$  such that  $t$  and its children take more than  $q$  steps, *excluding* synchronisation (that is, calls to **fork**, **newch**, **signal** and **wait**). Note that by Lemma 21,  $\mathcal{K}$  is signal-ordered.

Now we rearrange  $\mathcal{K}$  to give a sequentialised trace of length greater than  $q$ .

- By applying Lemma 28, we extend the trace until the first  $n$  children of  $t$  together take more than  $q$  non-synchronisation steps, and all children but the  $n$ th have terminated.
- By applying Lemma 29, we erase the  $(n+1)$ th to maximum child threads of  $t$ , along with associated calls to fork from  $t$ .

The resulting trace  $\mathcal{K}$  has more than  $q$  non-synchronisation transitions from  $t$  and its first  $n$  children, and in the final state  $\kappa'$ , all but the  $n$ th child thread have terminated. By applying Lemma 20, we can construct a trace  $\mathcal{K}''$  that has equivalent thread-local behaviour, and that is sequentialised with respect to  $t$ .

We observe that, by Lemma 22,  $C = \llbracket \llbracket C_{\text{par}} \rrbracket_{\text{fk}}^{\downarrow} \rrbracket_{\text{ch}}^{\downarrow}$ . By applying Lemmas 14 and 15, construct a trace  $\mathcal{K}_{\text{seq}}$  with initial state  $\kappa_{\text{seq}}$ . The invariants in Lemmas 14 and 15 only erase synchronisation transitions. Consequently, the resulting sequential trace has more than  $q$  transitions over the thread  $t$ . This contradicts our assumption, and completes the proof.  $\square$

## C Fully-tabulated Intermediate Assertions

Fig. C and Fig C show the needed and redundant maps computed by the resource-usage analysis for the example discussed in §5.1.

	$l_2$	$l_2l_3$	$l_2l_3l_5$	$l_2l_3g_e$	$l_2l_4$	$l_2l_4l_5$	$l_2l_4g_e$	$f_e$
$l_1$	$x \mapsto x$	$x \geq i \wedge x \mapsto x$	$x \geq i \wedge x \mapsto x$	$x \geq i \wedge x \mapsto x * y \mapsto y$	$x < i \wedge x \mapsto x$	$x < i \wedge x \mapsto x$	$x < i \wedge x \mapsto x$	$(x \geq i \wedge x \mapsto x * y \mapsto y) \vee (x < i \wedge x \mapsto x)$
$l_2$		$v = x \wedge v \geq i$	$v = x \wedge v \geq i$	$v = x \wedge v \geq i * y \mapsto y$	$v = x \wedge v < i$	$v = x \wedge v < i$	$v = x \wedge v < i \wedge x \mapsto x$	$(v = x \wedge v \geq i * y \mapsto y) \vee (v = x \wedge v < i \wedge x \mapsto x)$
$l_2l_3$			$v = x \wedge v \geq i$	$v = x \wedge v \geq i * y \mapsto y$				$v = x \wedge v \geq i * y \mapsto y$
$l_2l_3l_5$				$v = x \wedge v \geq i * y \mapsto y$				$v = x \wedge v \geq i * y \mapsto y$
$l_2l_3g_e$								$v = x \wedge v \geq i$
$l_2l_4$						$v = x \wedge v < i$	$v = x \wedge v < i * x \mapsto x$	$v = x \wedge v < i * x \mapsto x$
$l_2l_4l_5$							$v = x \wedge v < i * x \mapsto x$	$v = x \wedge v < i * x \mapsto x$
$l_2l_4g_e$								$v = x \wedge v < i$

Table 2: needed map computed by the algorithm NEEDED. (not shown trivial entries; empty entry means not defined)



	$l_2$	$l_2l_3$	$l_2l_3l_5$	$l_2l_3g_e$	$l_2l_4$	$l_2l_4l_5$	$l_2l_4g_e$	$f_e$
$l_1$	$y \mapsto y$	$(x \geq i \wedge y \mapsto y) \vee (x < i \wedge x \mapsto x * y \mapsto y)$	$(x \geq i \wedge y \mapsto y) \vee (x < i \wedge x \mapsto x * y \mapsto y)$	$x < i \wedge x \mapsto x * y \mapsto y$	$(x \geq i \wedge x \mapsto x * y \mapsto y) \vee (x < i \wedge y \mapsto y)$	$(x \geq i \wedge x \mapsto x * y \mapsto y) \vee (x < i \wedge y \mapsto y)$	$(x \geq i \wedge x \mapsto x * y \mapsto y) \vee (x < i \wedge y \mapsto y)$	$x < i \wedge y \mapsto y$
$l_2$		$v = x \wedge x \mapsto x * y \mapsto y$	$v = x \wedge x \mapsto x * y \mapsto y$	$(v = x \wedge v \geq i \wedge x \mapsto x) \vee (v = x \wedge v < i \wedge x \mapsto x * y \mapsto y)$	$v = x \wedge x \mapsto x * y \mapsto y$	$v = x \wedge x \mapsto x * y \mapsto y$	$(v = x \wedge v \geq i \wedge x \mapsto x * y \mapsto y) \vee (v = x \wedge v < i \wedge y \mapsto y)$	$(v = x \wedge v \geq i \wedge x \mapsto x) \vee (v = x \wedge v < i \wedge y \mapsto y)$
$l_2l_3$			$v = x \wedge v \geq i \wedge x \mapsto x * y \mapsto y$	$v = x \wedge v \geq i \wedge x \mapsto x$				$v = x \wedge v \geq i \wedge x \mapsto x$
$l_2l_3l_5$				$v = x \wedge v \geq i \wedge x \mapsto x$				$v = x \wedge v \geq i \wedge x \mapsto x$
$l_2l_3g_e$								$v = x \wedge v \geq i \wedge x \mapsto x * y \mapsto y$
$l_2l_4$						$v = x \wedge v < i \wedge x \mapsto x * y \mapsto y$	$v = x \wedge v < i \wedge y \mapsto y$	$v = x \wedge v < i \wedge y \mapsto y$
$l_2l_4l_5$							$v = x \wedge v < i \wedge y \mapsto y$	$v = x \wedge v < i \wedge y \mapsto y$
$l_2l_4g_e$								$v = x \wedge v < i \wedge x \mapsto x * y \mapsto y$

Table 3: redundant map computed by the algorithm REDUNDANT. (not shown trivial entries; empty entry means not defined)