

# Interleaving and Lock-Step Semantics for Analysis and Verification of GPU Kernels<sup>\*</sup>

Peter Collingbourne<sup>1\*\*</sup> Alastair F. Donaldson<sup>1</sup> Jeroen Ketema<sup>1</sup> Shaz Qadeer<sup>2</sup>

<sup>1</sup> Imperial College London

`peter@pcc.me.uk, {afd, jketema}@imperial.ac.uk`

<sup>2</sup> Microsoft Research

`qadeer@microsoft.com`

**Abstract.** Graphics Processing Units (GPUs) from leading vendors employ predicated (or guarded) execution to eliminate branching and increase performance. Similarly, a recent GPU verification technique uses predication to reduce verification of GPU kernels (the massively parallel programs that run on GPUs) to verification of a sequential program.

Prior work on the formal semantics of lock-step predicated execution for kernels focused on *structured programs*, where control is organised using **if**- and **while**-statements. We provide lock-step execution semantics for GPU kernels that are represented by *arbitrary* reducible control flow graphs. We present a traditional interleaving semantics and a novel lock-step semantics based on predication, and show that for terminating kernels either both semantics compute identical results or both behave erroneously.

The method allows reducing GPU kernel verification to the verification of a sequential, lock-step program to be applied to GPU kernels with arbitrary reducible control flow. We have implemented the method in the GPUVerify tool, and present an evaluation using a set of 163 open source and commercial GPU kernels. Among these kernels, 42 exhibit unstructured control flow which our novel lock-step predication technique can handle fully automatically. This generality comes at a modest price: verification across our benchmark set was on average 2.25 times slower than using an existing approach that specifically targets structured kernels.

## 1 Introduction

Graphics Processing Units (GPUs) have recently found application in accelerating general-purpose computations, e.g. in image retrieval [24] and machine learning [3]. If an application exhibits significant parallelism it may be possible to extract the computational core of the application as a *kernel* and offload this kernel to run across the parallel hardware of a GPU, sometimes beating CPU performance by orders of magnitude. Writing kernels for massively parallel GPUs is challenging, requiring coordination of a large number of threads. Data races and mis-synchronisation at barriers (known as *barrier divergence*) can lead to erroneous and non-deterministic program behaviours. Worse, they can lead to bugs which manifest *only* on some GPU architectures.

---

<sup>\*</sup> This work was supported by the EU FP7 STREP project CARP (project number 287767).

<sup>\*\*</sup> Peter Collingbourne is currently employed at Google.

Recently, substantial effort has been put into the design of tools for rigorous analysis of GPU kernels [7,17,8,18,16]. In prior work [7], we designed a verification technique and tool, GPUVerify, for OpenCL [14] and CUDA [22] kernels. GPUVerify achieves scalability by reducing verification of a parallel kernel to a *sequential* program verification task. This is achieved by transforming the kernel into a form where all threads execute in *lock-step* in a manner that still facilitates detection of data races arising due to arbitrary thread interleavings.

Semantics and program transformations for lock-step execution have been formally studied for *structured* GPU kernels where control flow is described by *if* and *while* constructs [7], but the technique of [7] is limited to the analysis of structured kernels. Lock-step semantics for GPU kernels where control flow is described by an *arbitrary* reducible control flow graph (CFG),<sup>3</sup> has *not* been studied. This is a serious limitation to the design of GPU kernel analysis techniques: kernels frequently exhibit unstructured control flow, either directly, e.g. through the use of *switch*-statements, or indirectly, through short-circuit evaluation of Boolean expressions. Dealing with CFGs also enables analysis of GPU kernels after compiler optimisations have been applied, bringing the analysis closer to the code actually executed by the GPU. It allows for the reuse of existing compiler infrastructures, such as Clang/LLVM, which use CFGs as their intermediate representation. Reusing compiler infrastructures hugely simplifies tool development, removing the burden of writing a robust front-end for C-like languages.

We present a traditional interleaving semantics and a novel lock-step semantics for GPU kernels described by CFGs. We show that if a GPU kernel is guaranteed to terminate then the kernel is correct with respect to the interleaving semantics if and only if it is correct with respect to the lock-step semantics, where *correct* means that all execution traces are free from data races, barrier divergence, and assertion failures. Our novel lock-step semantics enables the strategy of reducing verification of a multithreaded GPU kernel to verification of a sequential program to be applied to arbitrary GPU kernels, and we have implemented this method in the GPUVerify tool. We present an experimental evaluation, applying our new tool to a set of 163 open source and commercial GPU kernels. In 42 cases these kernels exhibited unstructured control flow either explicitly (e.g. through *switch*-statements), or implicitly due to short-circuit evaluation. In the former case, these kernels had to be manually simplified to be amenable to analysis using the original version of GPUVerify. In the latter case, it turned out that the semantics of short-circuit evaluation of logical operators was not handled correctly in GPUVerify. Our new, more general implementation handles all these kernels accurately and automatically. Our results show that overall GPUVerify continues to perform well, with verification across our benchmark set being on average only 2.25 times slower compared to the original version that was limited to structured kernels [7].

In summary, our main contributions are:

- A novel operational semantics for lock-step execution of GPU kernels with arbitrary reducible control flow.

---

<sup>3</sup> Henceforth, whenever we refer to a CFG we shall always mean a reducible CFG. For a definition of reducibility we refer the reader to [1]. We note that irreducibility is uncommon in practice. In particular, we have never encountered a GPU kernel with an irreducible control flow graph, and whether irreducible control flow is supported at all is implementation-defined in OpenCL [14].

<pre> __kernel void scan(__global int *sum) {   int offset = 1, temp;   while (offset &lt; TS) {     if (tid &gt;= offset)       temp = sum[tid - offset];     barrier();     if (tid &gt;= offset)       sum[tid] = sum[tid] + temp;     barrier();     offset *= 2;   } } </pre>	<pre> __kernel void scan(__global int *sum) {   int offset = 1, temp;   while (offset &lt;= tid) {     temp = sum[tid - offset];     barrier();     sum[tid] = sum[tid] + temp;     barrier();     offset *= 2;   } } </pre>
(a) A correct kernel	(b) A kernel with barrier divergence

Fig. 1: Two OpenCL kernels

- A proof-sketch that this semantics is equivalent to a traditional interleaving semantics for terminating GPU kernels.
- A verification method for GPU kernels described as CFGs, which uses our lock-step semantics to reduce verification of a multithreaded kernel to a sequential verification task, implemented in the GPUVerify tool.
- An experimental evaluation of this implementation with respect to a large collection of benchmarks.

After presenting a small example to provide some background on GPU kernels and illustrate the problems of data races and barrier divergence (Sect. 2), we present the interleaving semantics (Sect. 3), our novel lock-step semantics (Sect. 4) and a proof-sketch showing that the semantics are equivalent for terminating kernels (Sect. 5). We then discuss the implementation in GPUVerify, and present our experimental results (Sect. 6). We end with related work and conclusions (Sect. 7).

## 2 A Background Example

We use an example to illustrate the key concepts from GPU programming and provide an informal description as to how predicated lock-step execution works for structured programs. We return to this example when presenting interleaving and lock-step semantics for kernels described as CFGs in Sects. 3 and 4.

*Threads, Barriers, and Shared Memory.* Figure 1a shows an OpenCL kernel<sup>4</sup> to be executed by  $TS$  threads, where  $TS$  is a power of two. The kernel computes a *scan* operation on the `sum` array so that at the end of the kernel we have, for all  $0 \leq i < TS$ ,  $sum[i] = \sum_{j=0}^i \mathbf{old}(sum)[j]$ , where  $\mathbf{old}(sum)$  refers to the `sum` array at the start of the kernel. All threads execute this kernel function in parallel, and threads may follow different control paths or access distinct data by querying their unique thread id, `tid`.

<sup>4</sup> For ease of presentation we use a slightly simplified version of OpenCL syntax, and we assume that all threads reside in the same work group and that this work group is one dimensional. Our implementation, described in Sect. 6, supports OpenCL in full.

Communication is possible via shared memory; the `sum` array is marked as residing in global shared memory via the `_global` qualifier. Threads synchronise using a `barrier`-statement, a collective operation that requires all threads to reach the same syntactic barrier before any thread proceeds past the barrier.

*Data Races and Barrier Divergence.* Two common defects from which GPU kernels suffer are *data races* and *barrier divergence*. Accesses to `sum` inside the loop are guarded so that on loop iteration  $i$  only threads with `id` at least  $2^{i-1}$  access the `sum` array. If either of the barriers in the example were omitted the kernel would be prone to a *data race* arising due to thread  $t_1$  reading from `sum[t1 - offset]`, while thread  $t_2$  writes to `sum[t2]`, where  $t_2 = t_1 - \text{offset}$ . The kernel of Fig. 1b aims to optimise the original example by reducing branches inside the loop: threads are restricted to only execute the loop body if their `id` is sufficiently large. This optimisation is *erroneous*; given a barrier inside a loop, the OpenCL standard requires that either all threads or zero threads reach the barrier during a given loop iteration, otherwise *barrier divergence* occurs and behaviour is undefined. In Fig. 1b, thread 0 will not enter the loop at all and thus will *never* reach the first barrier, while all other threads will enter the loop and reach the barrier. Unfortunately, on mainstream GPU architectures from AMD and NVIDIA the kernel of Fig. 1b behaves identically to the kernel of Fig. 1a, meaning that this barrier divergence bug was not detected during testing. This is problematic because the erroneous kernel code is not portable across architectures which support OpenCL.

*Lock-Step Predicated Execution.* We informally describe lock-step execution for structured programs as used by GPUVerify [7] and which we here generalise to CFGs.

To achieve lock-step execution, GPUVerify transforms kernels into *predicated*

form [2]. The example of Fig. 2 illustrates the effect of applying predication to the example of Fig. 1a. A statement of the form  $e \Rightarrow \text{stmt}$  is a *predicated* statement which is a no-op if  $e$  is false, and has the same effect as  $\text{stmt}$  if  $e$  is true. Observe that the `if`-statements in the body of the loop have been predicated: the condition (which is the same for both statements) is evaluated into a Boolean variable  $q$ , the conditional statement is removed and the statements previously inside the conditional are predicated by the associated Boolean variable. Lock-step execution of the `while`-loop is achieved by evaluating the loop condition into a Boolean variable  $p$ , predicated all statements in the loop body by  $p$ , and recomputing  $p$  at the end of the loop body. The loop condition is replaced by a guard which evaluates to false if and only if the predicate variable  $p$  is false for *every* thread. Thus all threads continue to execute the loop until every thread is ready to leave the loop; when the loop condition becomes false for a given thread the thread simply performs no-ops during subsequent loop iterations.

In predicated form, the kernel does not exhibit any thread interleavings, and thus can be regarded as a *sequential*, vector program. GPUVerify exploits this fact to reduce

```

__kernel void
scan(__global int *sum) {
    bool p, q;
    int offset = 1, temp;
    p = (offset < TS);
    while (∃ t :: t.p) {
        q = (p && tid >= offset);
        q ⇒ temp = sum[tid - offset];
        p ⇒ barrier();
        q ⇒ sum[tid] = sum[tid] + temp;
        p ⇒ barrier();
        p ⇒ offset *= 2;
        p ⇒ p = (offset < TS);
    }
}

```

Fig. 2: Lock-step predicated execution for structured kernel of Fig. 1a

GPU kernel verification to a sequential program verification task. The full technique, described in [7], involves considering lock-step execution of an arbitrary *pair* of threads, rather than all threads, and employs instrumentation variables to efficiently check for data races that could manifest due to arbitrary thread interleavings.

Predication is easy to perform at the level of structured programs built using **if**- and **while**-statements. However, as discussed in the introduction, GPU kernels in general may exhibit unstructured control flow. In Sect. 4 we present a program transformation for predicated execution of GPU kernels that are described as CFGs.

### 3 Interleaving Semantics for GPU Kernels

We introduce a simple language for describing GPU kernels as CFGs, and present an interleaving semantics for this language. We do not include procedures in our presentation due to space reasons, however our results do extend to a procedural setting and procedures are supported by the implementation described in Sect. 6.

#### 3.1 Syntax

The syntax for our language is identical to the core of the Boogie programming language [5], except that it includes an additional **barrier** statement:

$$\begin{aligned} \text{Program} &::= \text{Block}^+ \\ \text{Block} &::= \text{BlockId} : \text{Stmts} \text{ goto } \text{BlockId}^+ ; \\ \text{Stmts} &::= \varepsilon \mid \text{Stmt} ; \text{Stmts} \\ \text{Stmt} &::= \text{Var} := \text{Expr} \mid \text{havoc } \text{Var} \mid \text{assume } \text{Expr} \mid \text{assert } \text{Expr} \mid \text{skip} \mid \text{barrier} \end{aligned}$$

Here,  $\varepsilon$  is an empty sequence of statements. The form of expressions is mostly irrelevant. We do assume presence of (a) *equality testing* ( $=$ ), (b) the standard Boolean operators, and (c) a ternary operator  $\text{Expr}_1 ? \text{Expr}_2 : \text{Expr}_3$ , which — like the operator from C — evaluates  $\text{Expr}_2$  in case  $\text{Expr}_1$  is *true* and evaluates  $\text{Expr}_3$  in case  $\text{Expr}_1$  is *false*.

To summarise in words, a kernel consists of a number of *basic blocks*, with each block consisting of a number of *statements* followed by a **goto** that non-deterministically chooses which block to execute next based on the provided *BlockIds*; non-deterministic choice in combination with **assumes** is used to model branches.

Because **gotos** only appear at the end of blocks there is a one-to-one correspondence between kernels and CFGs. We assume that all kernels have reducible CFGs, which means that cycles in a CFG are guaranteed to form *natural loops*. A natural loop has a unique *header* node, the single entry point to the loop, and one or more *back edges* going from a loop node to the header [1].

We assume that each kernel has at least a block labelled *Start*. This is the block from which execution of each thread commences. Moreover, no block is labelled *End*; instead the occurrence of *End* in a **goto** signifies that the program may terminate at this point.

Figure 3 shows the kernel of Fig. 1a encoded in our simple programming language. The example uses an array; we could easily include arrays in our formalisation of GPU kernel semantics but, for brevity, we do not.

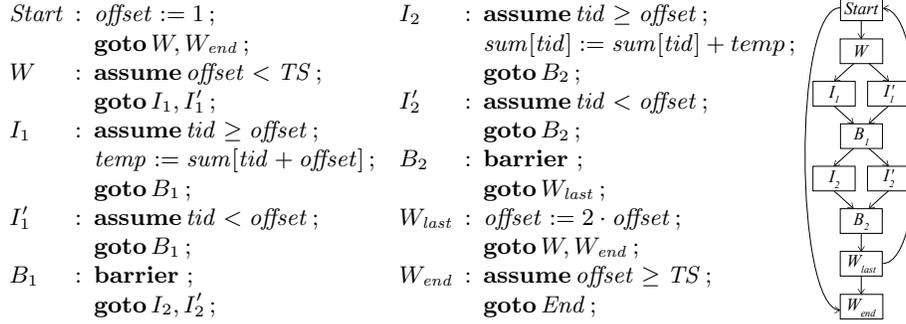


Fig. 3: The kernel of Fig. 1a encoded in our kernel language and its CFG

### 3.2 Operational Semantics

We now define a small-step operational semantics for our kernel programming language, which is based on interleaving the steps taken by individual threads.

*Individual Threads.* The behaviour of individual threads and non-barrier statements executed by these threads is presented in Figs. 4a and 4b.

The operational semantics of a thread  $t$  is defined in terms of triples  $\langle \sigma, \sigma_t, b_t \rangle$ , where  $\sigma$  is the GPU *shared* store,  $\sigma_t$  is the *private* store for thread  $t$  and  $b_t$  is the statement or sequence of statements the thread will execute, or more formally will *reduce*, next. Each store is a mapping from variables to the values in some domain  $D$ . We assume that *no* variable is mapped to a value in both  $\sigma$  and  $\sigma_t$ .

In Fig. 4a,  $(\sigma, \sigma_t)[v \mapsto val]$  denotes a pair of stores equal to  $(\sigma, \sigma_t)$  except that  $v$  (which we assume occurs in either  $\sigma$  or  $\sigma_t$ ) has been updated and is equal to  $val$ . In the same figure,  $(\sigma, \sigma_t)(e)$  denotes the evaluation of the expression  $e$  given  $(\sigma, \sigma_t)$ . The labels on the arrows allow us to observe (a) changes to stores and (b) the state of stores upon. A label is omitted when the stores do not change, e.g. in the SKIP-rule.

The symbols  $\surd$ ,  $\mathcal{E}$ , and  $\perp$  indicate, resp., *termination*, *error*, and *infeasible*. These are *termination statuses* and signify that a thread (or later kernel) has terminated with that particular status. Below, termination always means termination with status *termination*; termination with status *error* or *infeasible* is indicated explicitly.

The ASSIGN- and SKIP-rules of Fig. 4a are standard. The HAVOC-rule updates the value of a variable  $v$  with an arbitrary value from the domain  $D$  of  $v$ . The ASSERT<sub>T</sub> and ASSUME<sub>T</sub>-rules are no-ops if the assumption or assertion  $(\sigma, \sigma_t)(e)$  holds. If the assumption or assertion does not hold, ASSERT<sub>F</sub> and ASSUME<sub>F</sub> yield, resp.,  $\mathcal{E}$  and  $\perp$ .

In Fig. 4b,  $s$  denotes a statement and  $b$  denotes the *body of a block*, i.e. a sequence of statements followed by a **goto**. The SEQ<sub>B</sub>- and SEQ<sub>E,I</sub>-rules define reduction of  $s$ ;  $b$  in terms of reduction of  $s$ . The GOTO- and BLOCK-rules specify how reduction continues once the end of a block is reached. The END-rule specifies termination of a thread.

*Interleaving.* We give interleaving semantics for a kernel  $P$  with respect to a given thread count  $TS$  in Fig. 4c. This is defined in terms of tuples  $\langle \sigma, \langle \sigma_1, b_1 \rangle, \dots, \langle \sigma_{TS}, b_{TS} \rangle \rangle$ ,

$$\begin{array}{c}
\frac{val = (\sigma, \sigma_t)(e)}{P \vdash \langle \sigma, \sigma_t, v := e \rangle \xrightarrow{(\sigma, \sigma_t)} (\sigma, \sigma_t)[v \mapsto val]} \text{ ASSIGN} \\
\frac{val \in D}{P \vdash \langle \sigma, \sigma_t, \mathbf{havoc} v \rangle \xrightarrow{(\sigma, \sigma_t)} (\sigma, \sigma_t)[v \mapsto val]} \text{ HAVOC} \\
\frac{a \in \{\mathbf{assert}, \mathbf{assume}\} \quad (\sigma, \sigma_t)(e)}{P \vdash \langle \sigma, \sigma_t, a e \rangle \rightarrow (\sigma, \sigma_t)} \text{ ASSERT}_{\text{T}} \quad \frac{\neg(\sigma, \sigma_t)(e)}{P \vdash \langle \sigma, \sigma_t, \mathbf{assume} e \rangle \xrightarrow{(\sigma, \sigma_t)} \perp} \text{ ASSUME}_{\text{F}} \\
\frac{\neg(\sigma, \sigma_t)(e)}{P \vdash \langle \sigma, \sigma_t, \mathbf{assert} e \rangle \xrightarrow{(\sigma, \sigma_t)} \mathcal{E}} \text{ ASSERT}_{\text{F}} \quad \frac{}{P \vdash \langle \sigma, \sigma_t, \mathbf{skip} \rangle \rightarrow (\sigma, \sigma_t)} \text{ SKIP}
\end{array}$$

(a) Statement rules

$$\begin{array}{c}
\frac{P \vdash \langle \sigma, \sigma_t, s \rangle \xrightarrow{(\sigma, \sigma_t)} (\tau, \tau_t)}{P \vdash \langle \sigma, \sigma_t, s ; b \rangle \xrightarrow{(\sigma, \sigma_t)} \langle \tau, \tau_t, b \rangle} \text{ SEQ}_{\text{B}} \quad \frac{P \vdash \langle \sigma, \sigma_t, s \rangle \xrightarrow{(\sigma, \sigma_t)} e \quad e \in \{\mathcal{E}, \perp\}}{P \vdash \langle \sigma, \sigma_t, s ; b \rangle \xrightarrow{(\sigma, \sigma_t)} e} \text{ SEQ}_{\text{E,I}} \\
\frac{1 \leq i \leq n}{P \vdash \langle \sigma, \sigma_t, \mathbf{goto} B_1, \dots, B_n \rangle \rightarrow \langle \sigma, \sigma_t, B_i \rangle} \text{ GOTO} \quad \frac{(B : b) \in P}{P \vdash \langle \sigma, \sigma_t, B \rangle \rightarrow \langle \sigma, \sigma_t, b \rangle} \text{ BLOCK} \\
\frac{}{P \vdash \langle \sigma, \sigma_t, \mathbf{End} \rangle \xrightarrow{(\sigma, \sigma_t)} \checkmark} \text{ END}
\end{array}$$

(b) Thread rules

$$\begin{array}{c}
\frac{T_{\vec{\sigma}}|_t = \langle \sigma_t, b_t \rangle \quad P \vdash \langle \sigma, \sigma_t, b_t \rangle \xrightarrow{(\sigma, \sigma_t)} \langle \tau, \tau_t, c_t \rangle}{P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \xrightarrow{(\sigma, \vec{\sigma})} \langle \tau, T_{\vec{\sigma}}[\langle \tau_t, c_t \rangle]_t \rangle} \text{ THREAD}_{\text{B}} \\
\frac{T_{\vec{\sigma}}|_t = \langle \sigma_t, b_t \rangle \quad P \vdash \langle \sigma, \sigma_t, b_t \rangle \xrightarrow{(\sigma, \sigma_t)} \checkmark}{P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \xrightarrow{(\sigma, \vec{\sigma})} \langle \sigma, T_{\vec{\sigma}}[\langle \sigma_t, \checkmark \rangle]_t \rangle} \text{ THREAD}_{\text{T}} \\
\frac{T_{\vec{\sigma}}|_t = \langle \sigma_t, b_t \rangle \quad P \vdash \langle \sigma, \sigma_t, b_t \rangle \xrightarrow{(\sigma, \sigma_t)} s \quad s \in \{\mathcal{E}, \perp\}}{P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \xrightarrow{(\sigma, \vec{\sigma})} s} \text{ THREAD}_{\text{E,I}} \\
\frac{\forall 1 \leq t \leq TS : T_{\vec{\sigma}}|_t = \langle \sigma_t, \checkmark \rangle}{P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \xrightarrow{(\sigma, \vec{\sigma})} \checkmark} \text{ TERMINATION}
\end{array}$$

(c) Interleaving rules

$$\begin{array}{c}
\frac{T_{\vec{\sigma}}|_t = \langle \beta_t, \sigma_t, \mathbf{barrier} e_t ; b_t \rangle \wedge \neg(\sigma, \sigma_t)(e_t)}{P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \rightarrow \langle \sigma, T_{\vec{\sigma}}[\langle \beta_t, \sigma_t, b_t \rangle]_t \rangle} \text{ BARRIER}_{\text{SKIP}} \\
\frac{\forall t : T_{\vec{\sigma}}|_t = \langle \beta, \sigma_t, \mathbf{barrier} e_t ; b_t \rangle \wedge (\sigma, \sigma_t)(e_t)}{P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \rightarrow \langle \sigma, \langle \beta, \sigma_1, b_1 \rangle, \dots, \langle \beta, \sigma_{TS}, b_{TS} \rangle \rangle} \text{ BARRIER}_{\text{S}} \\
\frac{\forall t : T_{\vec{\sigma}}|_t = \langle \beta_t, \sigma_t, \mathbf{barrier} e_t ; b_t \rangle \wedge (\sigma, \sigma_t)(e_t) \quad \exists t_1, t_2 : T_{\vec{\sigma}}|_{t_1 \cdot \text{bv}} \neq T_{\vec{\sigma}}|_{t_2 \cdot \text{bv}}}{P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \xrightarrow{(\sigma, \vec{\sigma})} \mathcal{E}} \text{ BARRIER}_{\text{F}}
\end{array}$$

(d) Barrier synchronisation rules

Fig. 4: Interleaving operational semantics

where  $\sigma$  is the *shared store*,  $\sigma_t$  is the *private store* of thread  $t$ , and  $b_t$  is the program fragment thread  $t$  will reduce next. The private store of a thread is not accessible by any other thread. In the figure,  $T_{\vec{\sigma}}$  is shorthand for  $(\langle\sigma_1, b_1\rangle, \dots, \langle\sigma_{TS}, b_{TS}\rangle)$ , where  $\vec{\sigma} = (\sigma_1, \dots, \sigma_{TS})$ . Moreover,  $T_{\vec{\sigma}}|_t$  denotes  $\langle\sigma_t, b_t\rangle$  and  $T_{\vec{\sigma}}[\langle\sigma', b\rangle]_t$  denotes  $T_{\vec{\sigma}}$  with the  $t$ th element replaced by  $\langle\sigma', b\rangle$ .

The  $\text{THREAD}_B$ -rule defines how a single step is performed by a single thread, cf. the rules in Fig. 4b. The  $\text{THREAD}_T$ -rule defines termination of a single thread, where the thread enters the termination state  $\checkmark$  from which no further reduction is possible. The  $\text{THREAD}_{E,1}$ -rule specifies that a kernel should terminate with status *error* or *infeasible* in case one of the threads terminates as such. The  $\text{TERMINATION}$ -rule specifies that a kernel terminates when all threads have terminated.

The  $\text{THREAD}$  rules are non-deterministic and define an interleaving semantics, as a step might be possible in multiple threads. A thread *cannot* access the private store of any other thread, while the shared store is accessible by *all* threads.

Given stores  $\sigma, \sigma_1, \dots, \sigma_{TS}$ , we define a *reduction* of a kernel  $P$  with threads  $1 \leq t \leq TS$  as sequence of applications of the operational rules where each threads starts reduction from *Start* and where the *initial* shared store is  $\sigma$  and the *initial* private store of thread  $t$  is  $\sigma_t$ . A reduction is *maximal* if it is either infinite or if termination with status *termination*, *error*, or *infeasible* has occurred.

Our interleaving semantics effectively has a sequentially consistent memory model, which is not the case for GPUs in practice. However, because our viewpoint is that GPU kernels that exhibit data races should be regarded as erroneous, this is of no consequence.

*Barrier Synchronisation.* When we define lock-step predicated execution of barriers in Sect. 4 we will need to model execution of a barrier by a thread in a *disabled* state. In preparation for this, let us say that barrier statement has the form **barrier**  $e$ , with  $e$  a Boolean expression. In Sect. 4,  $e$  will evaluate to *true* if and only if the barrier is executed in an enabled state. The notion of thread-enabledness is not relevant to our interleaving semantics: we can view a thread as *always* being enabled. Thus we regard the **barrier** syntax of our kernel programming language as short for **barrier true**.

Figure 4d defines rules for (mis-)synchronisation between threads at barriers. Our aim here is to formalise the conditions for correct barrier synchronisation in OpenCL, which are stated informally in the OpenCL specification as follows [14]:

- B1 If **barrier** is inside a conditional statement, then all [threads] must enter the conditional if any [thread] enters the conditional statement and executes the barrier.
- B2 If **barrier** is inside a loop, all [threads] must execute the barrier for each iteration of the loop before any are allowed to continue execution beyond the barrier.

The rules of Fig. 4d capture these conditions using a number of special *barrier variables* that we assume are implicit in definition of each kernel:

- For every loop  $L$  in the kernel, every thread has a private *loop counter* variable  $v_L$ . The variable  $v_L$  of each thread  $t$  is initialised to zero, incremented each time the header node for  $L$  is reduced by  $t$ , and reset to zero on exit from  $L$ .
- Every thread has a private variable  $v_{\text{barrier}}$ . We assume that each barrier appearing in the kernel has a unique id. The variable  $v_{\text{barrier}}$  of each thread  $t$  is initialised to a

special value  $(-)$  different from every barrier id. When  $t$  reaches a barrier,  $v_{\text{barrier}}$  is set to the id of that barrier, and it is reset to  $(-)$  after reduction of the barrier.

The variable  $v_{\text{barrier}}$  codifies that each thread is synchronising on the same barrier, capturing condition  $B1$  above. The loop counters codify that each thread must have executed the same number of loop iterations upon synchronisation, capturing  $B2$ .

In Fig. 4d, we write  $(\beta_t, \sigma_t)$ , in case we want to make explicit the barrier variables  $\beta_t$  of each thread  $t$ . We write  $T_{\bar{\sigma}}|_{t.\text{bv}} = \beta_t$  where  $T_{\bar{\sigma}}|_t = \langle \beta_t, \sigma_t, b_t \rangle$ .

The  $\text{BARRIER}_{\text{SKIP}}$ -rule specifies that **barrier**  $e$  is a no-op if  $e$  is *false*. Although this can never occur for kernels written directly in our kernel programming language, our equivalence proof in Sect. 5 requires this detail to be accounted for.

The  $\text{BARRIER}_{\text{S}}$ -rule specifies that reduction continues beyond a barrier if all threads are at a barrier and the barrier variables agree across threads. The  $\text{BARRIER}_{\text{F}}$ -rule specifies that a kernel should terminate with *error* in case the threads have reached barriers with disagreeing barrier variables, i.e. when *barrier divergence* has occurred.

*Data Races.* We say that a thread  $t$  is *responsible* for a step in a reduction if a  $\text{THREAD}$  rule (see Fig. 4c) was employed in the step and the premise of the rule was instantiated with  $t$ . Moreover, we say that a thread  $t$  *accesses* a variable  $v$  in a step if  $t$  is responsible for the step and if in the step either (a) the value of  $v$  is used to evaluate an expression or (b)  $v$  is updated. The definition is now as follows:

**Definition 3.1.** *If  $P$  is a kernel, then  $P$  has data race if there is a maximal reduction  $\rho$  of  $P$  not ending in the infeasible status  $\perp$  and a shared variable  $v$  such that  $v$  is accessed by distinct threads  $t$  and  $t'$  during  $\rho$ , where at least one of the threads updates the variable and where no application of  $\text{BARRIER}_{\text{S}}$  occurs between accesses.*

*Terminating and Race Free Kernels.* We say that a kernel  $P$  is *terminating* with respect to the interleaving semantics if all maximal reductions of  $P$  are finite and do not end with status *error*. We say that  $P$  is *race free* with respect to the interleaving semantics if  $P$  has no data races according to Definition 3.1.

## 4 Lock-Step Semantics for GPU Kernels

We define lock-step execution semantics for GPU kernels represented as arbitrary CFGs in two stages. First, in Sect. 4.1, we present a program transformation which turns the sequential program executed by a single thread into a predicated form where control flow is *flattened*: all branches, except for loop back edges, are eliminated. Then, in Sect. 4.2, we use this transformation to express lock-step execution of all threads in a kernel as a sequential program in *vector* form. Each statement of this sequential program will perform the work of all threads in a single step.

To avoid many corner cases we assume that kernels always synchronise on a barrier immediately preceding termination. Moreover, if a block  $B$  ends with **goto**  $B_1, \dots, B_n$  then at most one of  $B_1, \dots, B_n$  is a loop head. A kernel can be trivially preprocessed to satisfy these restrictions.

*Sort Order.* Predication of CFGs involves flattening control flow, rewriting branches by predicating blocks and executing these blocks in a linear order. For CFGs without loops

any topological sort gives a suitable order: it ensures that if block  $B$  is a predecessor of  $C$  in the original CFG then  $B$  will appear before  $C$  in the predicated CFG. In the presence of loops the order must ensure that once execution of the blocks in a loop commences this loop will be executed completely before any node outside the loop is executed.

Formally, we require a total order  $\leq$  on blocks satisfying the following conditions:

- For all blocks  $B$  and  $C$ , if there is a path from  $B$  to  $C$  in the CFG, then  $B \leq C$  unless a back edge occurs on the path.
- For all loops  $L$ , if  $B \leq D$  and  $B, D \in L$ , then  $C \in L$  for all  $B \leq C \leq D$ .

A total order satisfying the above conditions always exists and can be computed as follows. Consider any innermost loop of the kernel and perform a topological sort of the blocks in the loop body (disregarding back edges). Replace the loop body by an abstract block. Repeat until no loops remain and perform a topological sort of resulting CFG. The sort order is now the order obtained by the final topological sort where we recursively replace each abstract node by the nodes it represents, i.e., if  $B \leq L \leq D$  with  $L$  an abstract node, then for any  $C \leq C'$  in the loop body represented by  $L$  we define  $B \leq C \leq C' \leq D$ . *End* can always be sorted last, as no **goto** occurs in it.

Considering the kernel of Fig. 3, we have that  $L = \{W, I_1, I'_1, B_1, I_2, I'_2, B_2, W_{last}\}$  is a loop and that  $Start \leq W \leq I_1 \leq I'_1 \leq B_1 \leq I_2 \leq I'_2 \leq B_2 \leq W_{last} \leq W_{end}$  satisfies our requirements; reversing  $I_1$  and  $I'_1$  or  $I_2$  and  $I'_2$  is possible.

In what follows we assume that a total order satisfying the above conditions has been chosen, and refer to this as the *sort order*. For a block  $B$  we use  $next(B)$  to denote the block that follows  $B$  in the sort order. If  $B$  is the final block in the sort order we define  $next(B)$  to be *End*, the block id denoting thread termination.

#### 4.1 Predication of a Single Thread

We now describe how predication of the body of a kernel thread is performed.

*Predication of Statements.* To predicate statements, we introduce a fresh private variable  $v_{active}$  for each thread, to which we assign *BlockIds*; the assigned *BlockId* indicates the block that needs to be executed.

If the value of  $v_{active}$  is not equal to the block that is currently being executed, all statements in the block will effectively be no-ops. In the case of **barrier** this follows by the **BARRIER<sub>SKIP</sub>**-rule.

Assuming the *BlockId* of the current block is  $B$ , predication of statements is defined in Table 1. In the case of **havoc**, the variable  $v_{havoc}$  is fresh and private.

Original form	Predicated form
$v := e;$	$v := (v_{active} = B) ? e : v;$
<b>havoc</b> $v;$	<b>havoc</b> $v_{havoc};$ $v := (v_{active} = B) ? v_{havoc} : v;$
<b>assert</b> $e;$	<b>assert</b> $(v_{active} = B) \Rightarrow e;$
<b>assume</b> $e;$	<b>assume</b> $(v_{active} = B) \Rightarrow e;$
<b>skip</b> ;	<b>skip</b> ;
<b>barrier</b> ;	<b>barrier</b> $(v_{active} = B);$

Table 1: Predication of statements

*Predication of Blocks.* Denoting by  $\pi(s)$  the predication of a single statement  $s$ , and applying  $\pi$  to sequences of statements in a pointwise fashion. Predication of blocks is now defined by default using the top row of Table 2, where  $v_{next}$  is a fresh, private

Original form	Predicated form
$B : ss$ <b>goto</b> $B_1, \dots, B_n$ ; <i>(B is not the last node of a loop according to the sort order)</i>	$B : \pi(ss)$ $v_{next} := \in \{B_1, \dots, B_n\}$ ; $v_{active} := (v_{active} = B) ? v_{next} : v_{active}$ ; <b>goto</b> $next(B)$ ;
$B : ss$ <b>goto</b> $B_1, \dots, B_n$ ; <i>(B is the last node of a loop according to the sort order)</i>	$B : \pi(ss)$ $v_{next} := \in \{B_1, \dots, B_n\}$ ; $v_{active} := (v_{active} = B) ? v_{next} : v_{active}$ ; <b>goto</b> $B_{back}, B_{exit}$ ; $B_{back} : \mathbf{assume} \ v_{active} = B_{head}$ ; <b>goto</b> $B_{head}$ ; $B_{exit} : \mathbf{assume} \ v_{active} \neq B_{head}$ ; <b>goto</b> $next(B)$ ;

Table 2: Predication blocks

variable, and  $:\in$  is shorthand for  $\mathbf{havoc} \ v_{next}$  ;  $\mathbf{assume} \ \bigvee_{i=1}^n (v_{next} = B_i)$ . Effectively,  $v_{active}$  is set to the value of the block that should be reduced next, while actual reduction will continue according to the sort order with block  $next(B)$ .

This method of predicating blocks does not deal correctly with loops: no block can be executed more than once as no back-edges are introduced. To take care of this, we predicate block  $B$  in a special manner if  $B$  belongs to a loop  $L$  and  $B$  occurs last in the sort order among all the blocks of  $L$ . Assume  $B_{head}$  is the head of  $L$ . The block  $B$  is predicated as in the bottom row of Table 2, where  $B_{back}$  and  $B_{exit}$  are fresh (see also Fig. 5). Our definition of a sort order guarantees that  $B_{head}$  is always sorted first among the blocks of  $L$ . By the introduction of  $B_{back}$ , reduction jumps back to  $B_{head}$  if  $L$  needs to be reduced again, otherwise reduction will continue beyond  $L$  by definition of  $B_{exit}$ .

```

Start : v_active := Start ;
      offset := (v_active = Start) ? 1 : offset ;
      v_next := {W, W_exit} ;
      v_active := (v_active = Start) ? v_next : v_active ;
      goto W ;
      :
B_2   : barrier (v_active = B_2) ;
      v_next := {W_last} ;
      v_active := (v_active = B_2) ? v_next : v_active ;
      goto W_last ;
W_last : offset :=
      (v_active = W_end) ? (2 · offset) : offset ;
      v_next := {W, W_exit} ;
      v_active := (v_active = W_last) ? v_next : v_active ;
      goto W_back, W_exit ;
W_back : assume v_active = F ;
      goto F ;
W_exit : assume v_active ≠ W ;
      goto W_end ;
W_end  : assume (v_active = W_end) ⇒ (offset ≥ TS) ;
      goto End ;

```

Fig. 5: Predication of the kernel of Fig. 3

*Predication of Kernels.* Predicating a complete kernel  $P$  now consists of three steps: (1) Compute a sort order on blocks as detailed above; (2) Predicate every block with

$$\begin{array}{c}
\frac{val = (\sigma, \sigma_t)(e)}{P \vdash \langle (\sigma, \sigma_t), v := e \rangle \xrightarrow{(\sigma, \sigma_t)} (\sigma, \sigma_t)[v \mapsto val]} \text{ ASSIGNS} \\
\\
\frac{val \in D}{P \vdash \langle (\sigma, \sigma_t), \mathbf{havoc} v \rangle \xrightarrow{(\sigma, \sigma_t)} (\sigma, \sigma_t)[v \mapsto val]} \text{ HAVOCS} \\
\\
\frac{\exists i : \sigma(e_i) \wedge val = (\sigma, \sigma_t)(e'_i)}{P \vdash \langle \sigma, v := \psi(\langle e_i, e'_i \rangle_{i=1}^n) \rangle \xrightarrow{(\sigma, \sigma_t)} \sigma[v \mapsto val]} \psi_T \\
\\
\frac{\forall i : \neg(\sigma, \sigma_t)(e_i)}{P \vdash \langle (\sigma, \sigma_t), v := \psi(\langle e_i, e'_i \rangle_{i=1}^n) \rangle \rightarrow (\sigma, \sigma_t)} \psi_F
\end{array}$$

Fig. 6: Operational semantics vector and synchronisation statements

respect to the sort order, according to the rules of Table 2; (3) Insert the assignment  $v_{\text{active}} := \text{Start}$  at the beginning of  $\pi(\text{Start})$ . The introduction of  $v_{\text{active}} := \text{Start}$  ensures that the statements from  $\pi(\text{Start})$  are always reduced first (see also Fig. 5).

## 4.2 Lock-Step Execution of All Threads

We now use the predication scheme of Sect. 4.1 for a single thread to define lock-step execution semantics for a kernel as a whole. We achieve this by encoding the kernel as a sequential program, each statement of which is a *vector* statement that performs the work of all threads simultaneously. To enable this, we first extend our programming language with these vector statements, as well as a statement related to barrier synchronisation.

*Vector and Synchronisation Statements.* We extend our language as follows:

$$\text{Stmt} ::= \dots \mid \text{Var}^* := \text{Expr}^* \mid \mathbf{havoc} \text{Var}^* \mid \text{Var} := \psi(\langle \text{Expr} \times \text{Expr} \rangle^*) \mid \mathbf{sync} \text{Expr}$$

The new vector assignment statement simultaneously assigns values to multiple variables, where we assume that the variables assigned to are all distinct and where the number of expressions is equal to the number of variables. Similarly, the new vector **havoc**-statement havocs multiple variables at once; again the variables are assumed to be distinct. The  $\psi$ -assignment will be used to model simultaneous writes to a shared variable by all threads. It takes a sequence  $(e_1, e'_1), \dots, (e_n, e'_n)$  with each  $e_i$  Boolean and non-deterministically assigns a value from  $\{\sigma(e'_i) \mid 1 \leq i \leq n \wedge \sigma(e_i)\}$  to the variable  $v$  (if the set is empty,  $v$  is left unchanged).

The **sync**-statement, behaves exactly as an **assert**; we introduce an additional keyword for assertions to be able to differentiate between assertions in our lock-step programs that originate, resp., from barriers and assertions.

The semantics for the new statements is presented in Fig. 6, where  $\langle e_i, e'_i \rangle_{i=1}^n$  denotes  $(e_1, e'_1), \dots, (e_n, e'_n)$ . Since **sync** behaves like **assert**, it is omitted from the figure.

*Lock-Step Execution.* To encode a kernel  $P$  as a single-threaded program  $\phi(P)$  which effectively executes all threads in lock-step, we assume for each *private* variable  $v$

Predicated form	Lock-step form
$v_{\text{active}} := \text{Start};$	$\langle v_{\text{active}} \rangle_{t=1}^{TS} := \langle \text{Start} \rangle_{t=1}^{TS}$
$v := (v_{\text{active}} = B) ? e : v;$	$v$ private $\langle v_t \rangle_{t=1}^{TS} := \langle (v_{\text{active},t} = B) ? \phi_t(e) : v_t \rangle_{t=1}^{TS};$ $v$ shared $v := \psi(\langle v_{\text{active},t} = B, \phi_t(e) \rangle_{t=1}^{TS});$
<b>havoc</b> $v_{\text{havoc}};$ $v := (v_{\text{active}} = B) ? v_{\text{havoc}} : v;$	$v$ private <b>havoc</b> $\langle v_{\text{havoc},t} \rangle_{t=1}^{TS};$ $\langle v_t \rangle_{t=1}^{TS} := \langle (v_{\text{active},t} = B) ? v_{\text{havoc},t} : v_t \rangle_{t=1}^{TS};$ $v$ shared <b>havoc</b> $\langle v_{\text{havoc},t} \rangle_{t=1}^{TS};$ $v := \psi(\langle v_{\text{active},t} = B, v_{\text{havoc},t} \rangle_{t=1}^{TS});$
<b>assert</b> $(v_{\text{active}} = B) \Rightarrow e;$	<b>assert</b> $\bigwedge_{t=1}^{TS} ((v_{\text{active},t} = B) \Rightarrow \phi_t(e))$
<b>assume</b> $(v_{\text{active}} = B) \Rightarrow e;$	<b>assume</b> $\bigwedge_{t=1}^{TS} ((v_{\text{active},t} = B) \Rightarrow \phi_t(e))$
<b>skip</b> ;	<b>skip</b> ;
<b>barrier</b> $(v_{\text{active}} = B);$	<b>sync</b> $(\bigvee_{t=1}^{TS} (v_{\text{active},t} = B)) \Rightarrow (\bigwedge_{t=1}^{TS} (v_{\text{active},t} = B));$

(a) Statements

Predicated form	Lock-step form
$B$ : $ss$ $v_{\text{next}} := \{B_1, \dots, B_n\};$ $v_{\text{active}} :=$ $(v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}};$ <b>goto</b> $next(B);$	$B$ : $\phi(ss)$ $\langle v_{\text{next},t} \rangle_{t=1}^{TS} := \langle \{B_1, \dots, B_n\} \rangle_{t=1}^{TS};$ $\langle v_{\text{active},t} \rangle_{t=1}^{TS} :=$ $\langle (v_{\text{active},t} = B) ? v_{\text{next},t} : v_{\text{active},t} \rangle_{t=1}^{TS};$ <b>goto</b> $next(B);$
$B$ : $ss$ $v_{\text{next}} := \{B_1, \dots, B_n\};$ $v_{\text{active}} :=$ $(v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}};$ <b>goto</b> $B_{\text{back}}, B_{\text{exit}};$ $B_{\text{back}}$ : <b>assume</b> $v_{\text{active}} = B_{\text{head}};$ <b>goto</b> $B_{\text{head}};$ $B_{\text{exit}}$ : <b>assume</b> $v_{\text{active}} \neq B_{\text{head}};$ <b>goto</b> $next(B);$	$B$ : $\phi(ss)$ $\langle v_{\text{next},t} \rangle_{t=1}^{TS} := \langle \{B_1, \dots, B_n\} \rangle_{t=1}^{TS};$ $\langle v_{\text{active},t} \rangle_{t=1}^{TS} :=$ $\langle (v_{\text{active},t} = B) ? v_{\text{next},t} : v_{\text{active},t} \rangle_{t=1}^{TS};$ <b>goto</b> $B_{\text{back}}, B_{\text{exit}};$ $B_{\text{back}}$ : <b>assume</b> $\bigvee_{t=1}^{TS} (v_{\text{active},t} = B_{\text{head}});$ <b>goto</b> $B_{\text{head}};$ $B_{\text{exit}}$ : <b>assume</b> $\bigwedge_{t=1}^{TS} (v_{\text{active},t} \neq B_{\text{head}});$ <b>goto</b> $next(B);$

(b) Blocks

Table 3: Lock-step construction

from  $P$  that there exists a variable  $v_t$  in  $\phi(P)$  for every  $1 \leq t \leq TS$ . For each *shared* variable  $v$  from  $P$  we assume there exists an identical variable in  $\phi(P)$ . Construction of a lock-step program for  $P$  starts from  $\pi(P)$  — the predicated version of  $P$ .

*Statements.* The construction for the predicated statements from Table 1 is presented in Table 3a. The construction involves making a copy of each statement for each thread. In the table,  $\phi_t$  denotes a map over expressions which replaces each private variable  $v$  by  $v_t$ . Note that for every thread  $t$ , there exists a variable  $v_{\text{active},t}$ , as variables freshly introduced by the predication scheme of Sect. 4.1 are private. Hence, we always know for each thread which block should be reduced next. We discuss the statements in turn.

Initially, each  $v_{\text{active},t}$  is assigned to *Start*. For every other assignment, we distinguish between assignments to private and shared variables. For a private variable  $v$ , the assignment is replaced by a vector assignment to the variables  $v_t$ , where  $\phi_t$  is applied to  $e$

$$\begin{aligned}
B_2 & : \mathbf{sync} \left( \bigvee_{t=1}^{TS} (v_{\text{active},t} = B) \right) \Rightarrow \left( \bigwedge_{t=1}^{TS} (v_{\text{active},t} = B) \right); \\
& \quad \langle v_{\text{next},t} \rangle_{t=1}^{TS} : \in \langle \{W_{\text{last}}\} \rangle_{t=1}^{TS}; \\
& \quad \langle v_{\text{active},t} \rangle_{t=1}^{TS} := \langle (v_{\text{active},t} = B_2) ? v_{\text{next},t} : v_{\text{active},t} \rangle_{t=1}^{TS}; \\
& \quad \mathbf{goto} W_{\text{last}}; \\
W_{\text{last}} & : \langle \text{offset}_t \rangle_{t=1}^{TS} := \langle (v_{\text{active},t} = F_{\text{end}}) ? (2 \cdot \text{offset}_t) : \text{offset}_t \rangle_{t=1}^{TS}; \\
& \quad \langle v_{\text{next},t} \rangle_{t=1}^{TS} : \in \langle \{F, F_{\text{exit}}\} \rangle_{t=1}^{TS}; \\
& \quad \langle v_{\text{active},t} \rangle_{t=1}^{TS} := \langle (v_{\text{active},t} = W_{\text{end}}) ? v_{\text{next},t} : v_{\text{active},t} \rangle_{t=1}^{TS}; \\
& \quad \mathbf{goto} W_{\text{back}}, W_{\text{exit}}; \\
W_{\text{back}} & : \mathbf{assume} \bigvee_{t=1}^{TS} (v_{\text{active},t} = W); \\
& \quad \mathbf{goto} W; \\
W_{\text{exit}} & : \mathbf{assume} \bigwedge_{t=1}^{TS} (v_{\text{active},t} \neq W); \\
& \quad \mathbf{goto} W_{\text{end}};
\end{aligned}$$

Fig. 7: Part of the lock-step program for the kernel of Fig. 3

as appropriate. For a shared variable  $v$ , it is not obvious which value needs to be assigned to  $v$ , as there might be multiple threads  $t$  with  $v_{\text{active},t} = B$ ; we non-deterministically pick the value from one of the threads with  $v_{\text{active},t} = B$ , employing a  $\psi$ -assignment.

In the case of a **havoc** followed by an assignment, there is again a case distinction between private and shared variables. For a private variable, the **havoc** and assignment are simply replaced by corresponding vector statements. For a shared variable, a vector havoc is used to produce an arbitrary value for each thread, and then the value associated with one of the threads  $t$  with  $v_{\text{active},t} = B$  non-deterministically assigned employing  $\psi$ .

In the case of **assert** and **assume**, we test whether  $(v_{\text{active},t} = B) \Rightarrow \phi_t(e)$  for each thread  $1 \leq t \leq TS$ . The **skip**-statement remains a no-op.

Lock-step execution of a **barrier** statement with condition  $v_{\text{active}} = B$  translates to an assertion checking that if  $v_{\text{active},t} = B$  holds for *some* thread  $t$  then it must hold for *all* threads. We shall sketch in Sect. 5 that checking for barrier divergence in this manner is equivalent to checking for barrier divergence using the interleaving semantics of Sect. 3. However, contrary to the interleaving case, there is no need to consider barrier variables in the lock-step case.

*Blocks.* The lock-step construction for blocks is presented in Table 3b, where  $\phi(ss)$  denotes the lock-step form for a sequence of statements.

If a block is *not* sorted last among the blocks of a loop (see the top row of Table 3b), we simply vectorise the updating  $v_{\text{active}}$ , where  $:\in$  is extended in the obvious way to non-deterministically assign values from multiple sets to multiple variables.

If a block is sorted last among blocks in a loops  $L$  then the successors of the block in the predicated program are  $B_{\text{back}}$ , which leads to the loop head, and  $B_{\text{exit}}$ , which leads to a node outside the loop. Our goal is to enforce the rule that *no* thread should leave the loop until *all* threads are ready to leave the loop, discussed informally in Sect. 2 and illustrated for structured programs by the guard of the **while** loop in Fig. 2. To achieve this, the bottom row of Table 3b uses an **assume** statement in  $B_{\text{back}}$  requiring that  $v_{\text{active}} = B_{\text{head}}$  for *some* thread, and an **assume** statement in  $B_{\text{exit}}$  requiring that  $v_{\text{active}} \neq B_{\text{head}}$  for *all* threads (see also Fig. 7 for a concrete example).

*Lock-Step Semantics and Data Races.* Having completed our definition of the lock-step construction  $\phi(P)$  for a kernel  $P$ , we now say that the lock-step semantics for  $P$  is the *interleaving semantics* for  $\phi(P)$ , with respect to a *single* thread (i.e. with  $TS = 1$ ). Our use of **sync** statements, which behave like assertions, captures the notion of barrier divergence. However, we need to define how data races can be detected by examining lock-step execution traces.

We say that thread  $t$  is *enabled* during a reduction step if the statement being reduced occurs in block  $B$  and  $v_{\text{active},t} = B$  holds at the point of reduction.

Let  $v$  be variable. We say that thread  $t$  *reads from*  $v$  during a reduction step if  $t$  is enabled during the reduction step and the reduction step involves evaluating an expression containing  $v$ . We say that  $t$  *writes to*  $v$  during a reduction step if  $t$  is enabled during the reduction step and the statement being reduced is an assignment to  $v$ . Note that in the case of a write, if multiple threads are enabled then  $v$  will be updated non-deterministically using one of the values supplied by the enabled threads. Nevertheless, we regard *all* enabled threads as having written to  $v$ .

A data race in a lock-step program is defined as follows:

**Definition 4.1.** *Let  $\phi(P)$  be the lock-step form of a kernel  $P$ . Then  $\phi(P)$  has a data race if there is a maximal reduction  $\rho$ , distinct threads  $t$  and  $t'$ , and a shared variable  $v$  such that:  $\rho$  does not end in *infeasible*;  $t$  writes to  $v$  during  $\rho$ ;  $t'$  either writes to or reads from  $v$  during  $\rho$ ; no **SYNC** occurs between the accesses (i.e. no barrier separates them).*

*Terminating and Race Free Kernels.* We say that a kernel  $P$  is *terminating* with respect to the lock-step semantics if all maximal reductions of  $\phi(P)$  are finite and do not end with status *error*. We say that  $P$  is *race free* with respect to the lock-step semantics if  $\phi(P)$  has no data races according to Definition 4.1.

## 5 Equivalence Between Interleaving and Lock-Step Semantics

We are now in a position to prove our main result, an equivalence between the kernels with interleaving semantics of Sect. 3 and lock-step programs of Sect. 4. The result applies to *well-formed* kernels, which we define as follows:

**Definition 5.1.** *A kernel  $P$  is well-formed if the following conditions hold for every block  $B$  in  $P$ :*

1. *The first statement of  $B$  is **assume**  $e$  where  $e$  refers only to private variables.*
2. *No other **assume**-statements appear in  $B$ .*
3. *If  $B$  ends with **goto**  $B_1, \dots, B_n$ , and every  $B_i$  begins with **assume**  $e_i$ , then  $\bigvee_{i=1}^n e_i$  is a tautology.*

Well-formedness means that, with the interleaving semantics, the *infeasible* status can only result from non-deterministic branching in a **goto** statement leading to a failing **assume** statement, and that in this case there is always an alternative way in which the branch could be resolved that would *not* lead to *infeasible*.

We can trivially enforce condition 1 of Definition 5.1 by inserting statements of the form **assume true** where necessary and removing shared variable accesses from

assumes via private temporary variables. Conditions 2 and 3 are guaranteed to hold if the CFG for  $P$  has been obtained from a kernel written in a C-like language such as OpenCL or CUDA.

Our main theorem, the following soundness and completeness result, applies to well-formed kernels:

**Theorem 5.2.** *Let  $P$  be a well-formed kernel and let  $\phi(P)$  be the lock-step version of  $P$ . Then,  $P$  is race free and terminating with respect to the interleaving semantics iff  $P$  is race free and terminating with respect to the lock-step semantics. Moreover, if race-freedom holds then for every terminating reduction of  $P$  there exists a terminating reduction of  $\phi(P)$ , and vice versa, such that every shared variable  $v$  has the same value at the end of both reductions.*

To see why well-formedness is required, consider the following kernel, where each thread  $t$  has a private variable  $tid$  whose value is  $t$  and where  $v$  is a shared:

```

Start : goto B1, B2;      B1 : assume tid = 1;      B2 : assume tid ≠ 1;
                                v := 4;                          v := 5;
                                assume v = 5;                       goto End;
                                goto End;

```

The interleaving semantics allows for reduction of `assume  $v = 5$`  after all assignments in all threads have taken place. Hence, if the assignment by thread 1 is not last among these,  $v = 5$  evaluates to `true`, and eventually termination occurs, with a data race. In the case of lock-step execution and assuming  $Start \leq B_1 \leq B_2$  in the sort order, we have that `assume  $v = 5$`  is always reduced immediately after  $v := 4$ . Hence, every reduction terminates with *infeasible* and no data race occurs.

That termination is required follows by adapting the counterexamples from [12,11] showing that CUDA hardware does not necessarily schedule threads from a non-terminating kernel in a way that that is fair from an interleaving point-of-view.

The proof of the theorem proceeds by showing that  $P$  and its predicated version  $\pi(P)$  are stutter equivalent, and then establishing a relationship between  $\pi(P)$  and  $\phi(P)$ .

*Equality of  $P$  and  $\pi(P)$ .* To show that  $P$  and  $\pi(P)$  are stutter equivalent [15], we define a denotational semantics of kernels is defined in terms of *execution traces* [5], i.e., sequences of tuples  $(\sigma, \vec{\sigma}) = (\sigma, \sigma_1, \dots, \sigma_{TS})$  with  $\sigma$  the shared store and  $\sigma_t$  the private store of thread  $t$ .

**Definition 5.3.** *Let  $\rho$  be a maximal reduction. The denotation or execution trace  $\mathcal{D}(\rho)$  of  $\rho$  is the sequence of  $\rightarrow$ -labels of  $\rho$  together with the termination status of  $\rho$  in case  $\rho$  terminates. Let  $(b_1, \dots, b_{TS})$  be a tuple block bodies or block ids. The denotation  $\mathcal{D}(b_1, \dots, b_{TS})$  of  $(b_1, \dots, b_{TS})$  is the set of denotations of all maximal reductions of  $(b_1, \dots, b_{TS})$  for all initial stores  $\sigma, \sigma_1, \dots, \sigma_t$  not terminating as *infeasible*. Let  $P$  be a kernel. The denotation  $\mathcal{D}(P)$  of  $P$  is  $\mathcal{D}(Start, \dots, Start)$ .*

Observe that *infeasible* traces are *not* included in the denotations of  $(b_1, \dots, b_{TS})$  and  $P$ ; these traces do not constitute actual program behaviour.

Stutter equivalence is defined on subsets of variables, where a *restriction* of a store  $\sigma$  to a set of variables  $V$  is denoted by  $\sigma \upharpoonright_V$  and, where given a tuple  $(\sigma, \vec{\sigma}) = (\sigma, \sigma_1, \dots, \sigma_{TS})$ , the restriction  $(\sigma, \vec{\sigma}) \upharpoonright_V$  is  $(\sigma \upharpoonright_V, \sigma_1 \upharpoonright_V, \dots, \sigma_{TS} \upharpoonright_V)$ .

**Definition 5.4.** Let  $V$  be a set of variables. Define the map  $\delta_V$  over execution traces as the map that replaces every maximal subsequence  $(\sigma_1, \vec{\sigma}_1) (\sigma_2, \vec{\sigma}_2) \cdots (\sigma_n, \vec{\sigma}_n) \cdots$  where  $(\sigma_1, \vec{\sigma}_1)|_V = (\sigma_2, \vec{\sigma}_2)|_V = \dots = (\sigma_n, \vec{\sigma}_n)|_V = \dots$  by  $(\sigma_1, \vec{\sigma}_1)$ .

Let  $\Sigma$  and  $T$  be execution traces. The traces are stutter equivalent with respect to  $V$ , denoted  $\Sigma \sim_{\text{st}}^V T$ , iff:

- $\Sigma$  and  $T$  are both finite with equal termination statuses and  $\delta_V(\Sigma) = \delta_V(T)$ ;
- $\Sigma$  and  $T$  are both infinite and  $\delta_V(\Sigma) = \delta_V(T)$ .

Let  $P$  and  $Q$  be kernels. The kernels are stutter equivalent with respect to  $V$ , denoted  $P \sim_{\text{st}}^V Q$ , iff for every  $\Sigma \in \mathcal{D}(P)$  there is a  $T \in \mathcal{D}(Q)$  with  $\Sigma \sim_{\text{st}}^V T$ , and vice versa.

Recall that  $\pi(P)$  denotes predication of  $P$ , we have the following:

**Theorem 5.5.** If  $P$  is a kernel with variables  $V$ , then  $\pi(P) \sim_{\text{st}}^V P$ . A data race occurs in  $P$  iff a data race occurs in  $\pi(P)$  where, during reduction of neither of the two statements causing the data race,  $v_{\text{active}} \neq B$  with  $B$  is the block containing the statement.

The result follows immediately by a case distinction on the statements that may occur in kernels once we establish the following lemma, which is a direct consequence of our construction and the first requirement on the sort order of blocks.

**Lemma 5.6.** Let  $P$  be a kernel with variables  $V$ . For any thread  $t$  and each block  $B$  of  $P$ , if  $(\sigma, \sigma_t)$  is a store of  $t$  and  $(\hat{\sigma}, \hat{\sigma}_t)$  is a store of in  $t$  in  $\pi(P)$  such that  $\hat{\sigma}|_V = \sigma$  and  $\hat{\sigma}(v_{\text{active},t}) = B$ , then

1. if the reduction of  $B$  is immediately followed by the reduction of a block  $C$ , then there exists a reduction of  $\pi(B)$  such that  $v_{\text{active},t} = C$  at the end of  $\pi(B)$  and eventually  $\pi(C)$  is reduced;
2. if the reduction of  $\pi(B)$  ends with  $v_{\text{active},t} = C$ , then there exists a reduction of  $B$  that is immediately followed by the reduction of a block  $C$ .

*Soundness and Completeness.* Theorem 5.2 is now proved as follows.

*Proof (Sketch<sup>5</sup>).* For termination and race-freeness of  $\phi(P)$ , it suffices by Lemma 5.6 to consider  $\pi(P)$  — the predicated version of  $P$ . Reason by contradiction and construct for a reduction of  $\phi(P)$  which is either infinite or has data race, a reduction of  $\pi(P)$  that also is either infinite or has a data race: Replace each statement and goto from the right-hand columns of Table 3 by a copy of the statement or goto in the left-hand column and reduce, where we introduce a copy for each thread. That SYNC can be replaced by BARRIER<sub>S</sub> follows as no statements from outside loops can be reduced while we are inside a loop (cf. the second requirement on sort order of blocks).

The remainder of the theorem follows by permuting steps of different threads so the reverse transformation from above can be applied.  $\square$

<sup>5</sup> A full proof of the theorem is provided at <http://multicore.doc.ic.ac.uk/tools/GPUVerify/ESOP2012/paper.pdf>.

## 6 Implementation and Experiments

*Implementation in GPUVerify.* We have implemented the predication technique described here in GPUVerify [7], a verification tool for OpenCL and CUDA kernels built on top of the Boogie verification engine [6] and Z3 SMT solver [20]. GPUVerify previously employed a predication technique for structured programs. Predication for CFGs has allowed us to build a new front-end for GPUVerify which takes LLVM intermediate representation (IR) as input; IR directly corresponds to a CFG. This allows us to compile OpenCL and CUDA kernels using the Clang/LLVM framework and perform analysis on the resulting IR. Hence, tricky syntactic features of C-like languages are taken care of by Clang/LLVM. Analysing kernels after compilation and optimisation also increases the probity of verification, opening up the opportunity to discover compiler-related bugs.

*Experimental Evaluation.* To assess the performance overhead in terms of verification time for our novel predication scheme and associated tool chain we compared our new implementation (GPUVerify II) with the original structured one (GPUVerify I).

We compared the tool versions using 163 OpenCL and CUDA kernels drawn from the AMD Accelerated Parallel Processing SDK v2.6 [4] (71 OpenCL kernels), the NVIDIA GPU Computing SDK v2.0 [21] (20 CUDA kernels), Microsoft C++ AMP Sample Projects [19] (20 kernels translated from C++ AMP to CUDA) and Rightware’s Basemark CL v1.1 [23] suite (52 OpenCL kernels, provided to us under an academic license). These kernels were used for analysis of GPUVerify I in [7], where several of the kernels had to be manually modified before they could be subjected to analysis: 4 kernels exhibited unstructured control flow due to `switch`-statements, and one featured a `do-while`-loop which was beyond the scope of the predication scheme of [7]. Furthermore, unstructured control flow arising from short-circuit evaluation of logical operators had been overlooked in GPUVerify I, which affected 30 of our example kernels. In GPUVerify II all kernels are handled uniformly due to our novel predication scheme in combination with the use of Clang/LLVM, which removes short-circuit evaluation in favour of unstructured control flow.

All experiments were performed on a PC with a 3.6 GHz Intel i5 CPU, 8 GB RAM running Windows 7 (64-bit), using Z3 v4.1. All times reported are averages over 3 runs. Both tool versions and all our benchmarks, except the commercial Basemark CL kernels, are available online to make our results reproducible:

<http://multicore.doc.ic.ac.uk/tools/GPUVerify/ESOP2012>

The majority of our benchmark kernels could be automatically verified by both GPUVerify I and GPUVerify II; 22 kernels were beyond the scope of both tools and result in a failed proof attempt. Key to the usability of GPUVerify is its *response time*, the time the tool takes to either report successful verification vs. a failed proof attempt. Comparing GPUVerify I and GPUVerify II we found that across the entire benchmark set the analysis time taken by GPUVerify II was 2.25 times that of GPUVerify I. The average and longest analysis time across all kernels were 4 seconds and 157 seconds respectively for GPUVerify I, and 10 seconds and 300 seconds respectively for GPUVerify II. Thus overall the accurate handling of unstructured control flow afforded by GPUVerify II comes at the price of a moderate performance penalty.

On 124 of the 163 kernels (76%), GPUVerify II was marginally (though not significantly) faster than GPUVerify I. For a further 21 kernels (13%) GPUVerify II was up to 50% slower than GPUVerify I. The remaining 18 kernels (11%) caused the slow down on average. In each case the time it took to run the front-ends and predication engines of both tool chains was negligible; the difference lay in constraint solving times; the SMT queries generated by our CFG-based tool chain can be somewhat more complex than in the structured case. The most dramatic example is a kernel which was verified by GPUVerify I and GPUVerify II in 3 and 202 seconds, resp., a slow-down for GPUVerify II of 70 times. This kernel exhibits a large number of shared memory accesses. In the LLVM IR processed by GPUVerify II these accesses are expressed as many separate, contiguous loads and stores, requiring reasoning about race-freedom between many pairs of operations. The structured approach of GPUVerify I captures these accesses at the abstract syntax tree level, allowing a load/store from/to a contiguous region to be expressed as a single access, significantly simplifying reasoning. This illustrates that there are benefits to working at the higher level of abstract syntax trees, and suggests that we might implement optimisations in GPUVerify II to automatically identify and merge contiguous memory accesses.

## 7 Related Work and Conclusion

*Related Work.* Interleaving semantics for GPU kernels has been defined by [16,18,12]. These are similar to our semantics, except that [16,12] do not give a semantics for barriers. Contrary to our lock-step approach, [16,18] battle the state space explosion due to arbitrary interleavings of threads by considering one particular schedule.

In [11,12], a semantics of CUDA kernels is defined that tries to model NVIDIA hardware as faithfully as possible. The focus is, however, not on predicated execution (although it does figure briefly in [11]), but on so-called *immediate post-dominator re-convergence* [10], a method to continue lock-step execution of threads as soon as possible after branch divergence has occurred between threads.

In addition to the above and similar to us, [12] shows for terminating kernels that CUDA execution of kernels can be faithfully simulated by certain interleaving thread schedules. The reverse is not shown; our analysis is that such a result is difficult to establish due to data races that occur in their examples.

*Conclusion.* Our lock-step semantics for GPU kernels expressed as arbitrary reducible CFGs enables automated analysis of a wider class of GPU kernels than previous techniques for structured programs, and allows for the analysis of compiled kernel code, after optimisations have been applied. Our soundness and completeness result establishes an equivalence between our lock-step semantics and a traditional semantics based on interleaving, and our implementation in GPUVerify and associated experimental evaluation demonstrate that our approach is practical.

Because our kernel programming language supports non-deterministic choice and hawking of variables it can express an over-approximation of a concrete kernel. In future work we plan to exploit this, investigating the combination of source-level abstraction techniques such as predicate abstraction with GPUVerify's verification method.

The well-formedness restrictions of Definition 5.1 mean that our equivalence result does not apply to kernels that exhibiting “dead end” paths. This is relevant if such paths are introduced through under-approximation, e.g., unwinding a loop by a fixed number of iterations in the style of bounded model checking. We plan to investigate whether it is possible to relax these well-formedness conditions under certain circumstances.

## References

1. Aho, A.V., Lam, M.S., Sethi, R., Ullman, J.D.: *Compilers: Principles, Techniques, and Tools*. Pearson Education, 2nd edn. (2007)
2. Allen, J., Kennedy, K., Porterfield, C., Warren, J.: Conversion of control dependence to data dependence. In: *POPL’83*. pp. 177–189 (1983)
3. Alshawabkeh, M., Jang, B., Kaeli, D.: Accelerating the local outlier factor algorithm on a GPU for intrusion detection systems. In: *GPGPU-3*. pp. 104–110 (2010)
4. AMD: AMD Accelerated Parallel Processing (APP) SDK, <http://developer.amd.com/sdks/amdappsdk/pages/default.aspx>
5. Barnett, M., Leino, K.R.M.: Weakest-precondition of unstructured programs. In: *PASTE’05*. pp. 82–87 (2005)
6. Barnett, M., et al.: Boogie: A modular reusable verifier for object-oriented programs. In: *FMCO 2005*. LNCS, vol. 4111, pp. 364–387 (2005)
7. Betts, A., Chong, N., Donaldson, A.F., Qadeer, S., Thomson, P.: GPUVerify: a verifier for GPU kernels. In: *OOPSLA 2012* (2012)
8. Collingbourne, P., Cadar, C., Kelly, P.H.J.: Symbolic testing of OpenCL code. In: *HVC 2011* (2012)
9. DeLine, R., Leino, K.R.M.: BoogiePL: A typed procedural language for checking object-oriented programs. Tech. Rep. MSR-TR-2005-70, Microsoft Research (2005)
10. Fung, W.W., Sham, I., Yuan, G., Aamodt, T.M.: Dynamic warp formation and scheduling for efficient GPU control flow. In: *MICRO 2007*. pp. 407–418 (2007)
11. Habermaier, A.: The model of computation of CUDA and its formal semantics. Tech. Rep. 2011-14, University of Augsburg (2011)
12. Habermaier, A., Knapp, A.: On the correctness of the SIMT execution model of GPUs. In: *ESOP 2012*. LNCS, vol. 7211, pp. 316–335 (2012)
13. Hecht, M.S., Ullman, J.D.: Characterizations of reducible flow graphs. *Journal of the ACM* 21(3), 367–375 (1974)
14. Khronos Group: The OpenCL specification, version 1.2 (2011)
15. Lamport, L.: What good is temporal logic? In: *Information Processing 83*. pp. 657–668 (1983)
16. Leung, A., Gupta, M., Agarwal, Y., Gupta, R., Jhala, R., Lerner, S.: Verifying GPU kernels by test amplification. In: *PLDI 2012*. pp. 383–394 (2012)
17. Li, G., Gopalakrishnan, G.: Scalable SMT-based verification of GPU kernel functions. In: *FSE 2010*. pp. 187–196 (2010)
18. Li, G., Li, P., Sawaya, G., Gopalakrishnan, G., Ghosh, I., Rajan, S.P.: GKLEE: concolic verification and test generation for GPUs. In: *PPoPP 2012*. pp. 215–224 (2012)
19. Microsoft Corporation: C++ AMP sample projects for download, <http://blogs.msdn.com/b/nativeconcurrency/archive/2012/01/30/c-amp-sample-projects-for-download.aspx>
20. de Moura, L.M., Bjørner, N.: Z3: An efficient SMT solver. In: *TACAS 2008*. LNCS, vol. 4963, pp. 337–340. Springer (2008)
21. NVIDIA: CUDA Toolkit Release Archive, <http://developer.nvidia.com/cuda/cuda-toolkit-archive>

22. NVIDIA: NVIDIA CUDA C Programming Guide, Version 4.2 (2012)
23. Rightware Oy: Basemark CL, <http://www.rightware.com/en/Benchmarking+Software/Basemark\%99+CL>
24. Zhu, F., Chen, P., Yang, D., Zhang, W., Chen, H., Zang, B.: A GPU-based high-throughput image retrieval algorithm. In: GPGPU-5. pp. 30–37 (2012)

## A Control Flow Graphs

The definition of *control flow graphs (CFGs)* is as usual [1]. Given a CFG  $G$ , a node  $d$  *dominates a node*  $n$ , if every path from the root of  $G$  to  $n$  goes through  $d$ . A *back edge* is an edge whose target dominates its source. A CFG  $G$  is *reducible* if the maximal acyclic CFG that is a subgraph of  $G$  is unique [13]. In a reducible CFG, the maximal direct acyclic subgraph is obtained by removing all back edges. A *loop with head*  $n$  in a reducible CFG is a maximal set of nodes and edges  $L$  with a node  $n \in L$  such that (a)  $n$  dominates all other nodes in  $L$  and (b) for every node  $m \in L$  there is a path to  $n$ . The *body* of a loop  $L$  is the set of all nodes and edges in  $L$  except for the back edges.

## B Omitted Proofs

To fully prove Theorem 5.2, we first consider the one-threaded, sequential language embedded in our kernel language, i.e. the language whose rules occur in Figs. 4b and 4a. Then, for full generality we extend this language with procedures. Having finished the discussion of the sequential language, we prove Theorem 5.5. This is followed by the actual proof of Theorem 5.2.

Recall that any newly introduced variables or block identifiers are called *fresh*.

### B.1 Stutter Equivalence of Sequential Programs

To allow us to prove results about statements or sequences of statements in isolation, without having to consider blocks, we introduce one additional rule on the level of statements:

$$\frac{P \vdash \langle \sigma, \sigma_t, s \rangle \xrightarrow{(\sigma, \sigma_t)} (\tau, \tau_t)}{P \vdash \langle \sigma, \sigma_t, s \rangle \xrightarrow{(\sigma, \sigma_t)} \langle \tau, \tau_t, \text{End} \rangle} \text{SEQs}$$

Thus, termination occurs after we reduce a statement in complete isolation.

For the purpose of this section, we assume that the above rule and the rules in Figs. 4b and 4a are defined over a *single store*, where we write  $\langle \sigma, b \rangle$  and where the  $\rightarrow$ -labels are simplified accordingly to  $\sigma$ . The proofs of this section carry over to the two-store setting, mutatis mutandis. We ignore the **barrier**-statement for the purposes of this section.

*Remark B.1.* Observe that  $\sigma[v \mapsto \sigma(v)] = \sigma$ , i.e. assigning a variable the value it already has does *not* modify the store. Observe also that the only rules that can modify a store are the ASSIGN- and HAVOC-rules, which only modify the value of the variable being assigned to or havocked; evaluation of expressions does *not* modify the store.

**Denotational Semantics.** The denotational semantics for the sequential language is defined as follows, cf. Definition 5.3:

**Definition B.2.** Let  $\rho$  be a maximal reduction. The denotation or execution trace  $\mathcal{D}(\rho)$  of  $\rho$  is the sequence of  $\rightarrow$ -labels of  $\rho$  together with the termination status of  $\rho$  in case  $\rho$  terminates. Let  $b$  be a statement, a sequence of statements, or a block. The denotation  $\mathcal{D}(b)$  of  $b$  is the set of denotations of all maximal reductions of  $b$  for all initial stores  $\sigma$  not terminating as infeasible. Let  $P$  be a program. The denotation  $\mathcal{D}(P)$  of  $P$  is  $\mathcal{D}(\text{Start})$ .

As before, that *infeasible* traces are *not* included in the denotations of  $b$  and  $P$ ; these traces do not constitute actual program behaviour.

The following is immediate by the observation that the END- and BLOCK-rules do not modify the store:

**Proposition B.3.** If  $ss$  and  $tt$  are sequences of statements, then  $\mathcal{D}(ss; tt)$  is the set of all infinite and error traces of  $ss$  together with all traces  $\Sigma \tau \mathcal{Y}$  where  $\Sigma \tau$  is a terminating trace of  $ss$  and  $\tau \mathcal{Y} \in \mathcal{D}(tt)$ .

If  $ss$  is a sequence of statements and  $B_1, \dots, B_n$  are blocks, then it holds that  $\mathcal{D}(ss \text{ goto } B_1, \dots, B_n;)$  is the set of all infinite and error traces of  $ss$  together with all traces  $\Sigma \tau \mathcal{Y}$  where  $\Sigma \tau$  is a terminating trace of  $ss$  and  $\tau \mathcal{Y} \in \mathcal{D}(B_i)$  for some  $1 \leq i \leq n$ .

As the SKIP-rule has an empty  $\rightarrow$ -label, we also have:

**Proposition B.4.** If  $ss$  is a sequence of statements, then

$$\mathcal{D}(ss) = \mathcal{D}(\text{skip}; ss) = \mathcal{D}(ss \text{ skip};).$$

As mentioned in Sect. 5, we compare programs based on their denotations by means of a stutter equivalence on subsets of variables. Recall that the *restriction* of a store  $\sigma$  to a set of variables  $V$  is denoted by  $\sigma \upharpoonright_V$ . Stutter equivalence for our sequential language is defined as follows, cf. Definition 5.4:

**Definition B.5.** Let  $V$  be a set of variables. Define the map  $\delta_V$  over execution traces as the map that replaces every maximal consecutive subsequence  $\sigma_1 \sigma_2 \dots \sigma_n \dots$  where  $\sigma_1 \upharpoonright_V = \sigma_2 \upharpoonright_V = \dots = \sigma_n \upharpoonright_V = \dots$  by  $\sigma_1 \upharpoonright_V$ .

Let  $\Sigma$  and  $T$  be execution traces. The traces are stutter equivalent with respect to  $V$ , denoted  $\Sigma \sim_{\text{st}}^V T$ , iff:

- $\Sigma$  and  $T$  are both finite with equal termination statuses and  $\delta_V(\Sigma) = \delta_V(T)$ ;
- $\Sigma$  and  $T$  are both infinite and  $\delta_V(\Sigma) = \delta_V(T)$ .

Let  $ss$  and  $tt$  be sequences of statements. The sequences are stutter equivalent with respect to  $V$ , denoted  $ss \sim_{\text{st}}^V tt$ , iff for every  $\Sigma \in \mathcal{D}(ss)$  there exists a  $T \in \mathcal{D}(tt)$  such that  $\Sigma \sim_{\text{st}}^V T$ , and vice versa.

Let  $B$  and  $C$  be blocks and let  $ss$  and  $tt$  be the statements of, resp.,  $B$  and  $C$ . The blocks are block level stutter equivalent with respect to  $V$ , denoted  $B \sim_{\text{st}}^V C$ , iff  $ss \sim_{\text{st}}^V tt$ .

Let  $P$  and  $Q$  be programs. The programs are stutter equivalent with respect to  $V$ , denoted  $P \sim_{\text{st}}^V Q$ , iff for every  $\Sigma \in \mathcal{D}(P)$  there exists a  $T \in \mathcal{D}(Q)$  such that  $\Sigma \sim_{\text{st}}^V T$ , and vice versa.

Observe that block level stutter equivalence is only concerned with the statements in blocks and *disregards* gotos.

**Stutter Equivalence of Statements.** Denote predication of statements by  $\pi$ , we have:

**Lemma B.6.** *Let  $P$  be a sequential program with variables  $V$ . For every statement  $s$  of  $P$  in a block  $B$*

1. *if  $v_{\text{active}} = B$ , then  $\pi(s) \sim_{\text{st}}^V s$ ;*
2. *if  $v_{\text{active}} \neq B$ , then  $\pi(s) \sim_{\text{st}}^V \text{skip}$ .*

*Proof.* We consider  $v_{\text{active}} = B$  and  $v_{\text{active}} \neq B$  in turn. Suppose  $v_{\text{active}} = B$ . Given a store  $\sigma$  of  $P$  and a store  $\hat{\sigma}$  of  $\pi(P)$  with  $\hat{\sigma}|_V = \sigma$ , we have:

- If  $s$  is  $v := e$ , then the only trace of  $s$  is the terminating trace  $\sigma \tau$  with  $\tau = \sigma[v \mapsto \sigma(e)]$ . As  $v_{\text{active}} = B$ , the only trace of  $\pi(s)$  is the terminating trace  $\hat{\sigma} \hat{\tau}$  with  $\hat{\tau} = \hat{\sigma}[v \mapsto \hat{\sigma}(e)]$ . Hence, as  $e$  only depends on variables in  $V$ , for  $e$  originates from  $P$ , we have  $\hat{\tau}|_V = \tau|_V$ .
- If  $s$  is **havoc**  $v$ , then the traces of  $s$  are the terminating traces  $\sigma v_{\text{val}}$  with  $v_{\text{val}} = \sigma[v \mapsto \text{val}]$  for  $\text{val} \in D$ . As  $v_{\text{active}} = B$ , the traces of  $\pi(s)$  are the terminating traces  $\hat{\sigma} \hat{\tau}_{\text{val}} \hat{v}_{\text{val}}$  with  $\hat{\tau}_{\text{val}} = \hat{\sigma}[v_{\text{havoc}} \mapsto \text{val}]$  and  $\hat{v}_{\text{val}} = \hat{\tau}_{\text{val}}[v \mapsto \text{val}]$ . As  $v_{\text{havoc}}$  is fresh,  $\hat{\sigma}|_V = \hat{\tau}_{\text{val}}|_V$  and  $\hat{v}_{\text{val}}|_V = \hat{\sigma}[v \mapsto \text{val}]|_V$ . Hence,  $\hat{\tau}_{\text{val}}|_V = \sigma|_V$  and  $\hat{v}_{\text{val}}|_V = v_{\text{val}}|_V$  for every  $\text{val} \in D$ .
- If  $s$  is **assert**  $e$ , then the only trace of  $s$  is (a) the terminating trace  $\sigma$  in case  $\sigma(e)$  holds, or (b) the *error* trace  $\sigma$  in case  $\neg\sigma(e)$  holds. As  $v_{\text{active}} = B$ , the only trace of  $\pi(s)$  is (a) the terminating trace  $\hat{\sigma}$  in case  $\hat{\sigma}(e)$  holds, and (b) the *error* trace  $\hat{\sigma}$  in case  $\neg\hat{\sigma}(e)$  holds. Hence, as  $e$  only depends on variables in  $V$ , the termination statuses of the traces  $\sigma$  and  $\hat{\sigma}$  are equal.
- If  $s$  is **assume**  $e$ , then the only trace of  $s$  is (a) the terminating trace  $\sigma$  in case  $\sigma(e)$  holds, or (b) the *infeasible* trace  $\sigma$  in case  $\neg\sigma(e)$  holds. As  $v_{\text{active}} = B$ , the only trace of  $\pi(s)$  is (a) the terminating trace  $\hat{\sigma}$  in case  $\hat{\sigma}(e)$  holds, and (b) the *infeasible* trace  $\hat{\sigma}$  in case  $\neg\hat{\sigma}(e)$  holds. Hence, as  $e$  only depends on variables in  $V$ , the termination statuses of the traces  $\sigma$  and  $\hat{\sigma}$  are equal.
- If  $s$  is **skip**, then the only trace of  $s$  is the terminating trace  $\sigma$  and the only trace of  $\pi(s)$  is the terminating trace  $\hat{\sigma}$ .

As  $\sigma$  is a store of  $P$ , we have  $\sigma|_V = \sigma$ . Thus, as  $\sigma$  and  $\hat{\sigma}$  are arbitrary with  $\hat{\sigma}|_V = \sigma$ , it follows by the above that  $\pi(s) \sim_{\text{st}}^V s$  in case  $v_{\text{active}} = B$ .

Suppose  $v_{\text{active}} \neq B$ . Given a store  $\sigma$  of  $P$ , the only trace of **skip** is the terminating trace  $\sigma$ . Moreover, given a store  $\hat{\sigma}$  of  $\pi(P)$  such that  $\hat{\sigma}|_V = \sigma$ :

- Suppose  $s$  is  $v := e$ . As  $v_{\text{active}} \neq B$ , the only trace of  $\pi(s)$  is the terminating trace  $\hat{\sigma} \hat{\tau}$  with  $\hat{\tau} = \hat{\sigma}[v \mapsto \hat{\sigma}(v)] = \hat{\sigma}$  (see also Remark B.1).
- Suppose  $s$  is **havoc**  $v$ . As  $v_{\text{active}} \neq B$ , the traces of  $\pi(s)$  are the terminating traces  $\hat{\sigma} \hat{\tau}_{\text{val}} \hat{v}_{\text{val}}$  with  $\hat{\tau}_{\text{val}} = \hat{\sigma}[v_{\text{havoc}} \mapsto \text{val}]$  and  $\hat{v}_{\text{val}} = \hat{\tau}_{\text{val}}[v \mapsto \hat{\tau}_{\text{val}}(v)] = \hat{\tau}_{\text{val}}$ . Hence, as  $v_{\text{havoc}}$  is fresh,  $\hat{\sigma}|_V = \hat{\tau}_{\text{val}}|_V = \hat{v}_{\text{val}}|_V$ .
- Suppose  $s$  is **assert**  $e$ . As  $v_{\text{active}} \neq B$ , the only trace of  $\pi(s)$  is the terminating trace  $\hat{\sigma}$ .

- Suppose  $s$  is **assume**  $e$ . As  $v_{\text{active}} \neq B$ , the only trace of  $\pi(s)$  is the terminating trace  $\hat{\sigma}$ .
- Suppose  $s$  is **skip**. The only terminating trace of  $\pi(s)$  is the terminating trace  $\hat{\sigma}$ .

As  $\sigma$  is a store of  $P$ , we have  $\sigma|_V = \sigma$ . Thus, as  $\sigma$  and  $\hat{\sigma}$  are arbitrary with  $\hat{\sigma}|_V = \sigma$ , it follows that  $\pi(s) \sim_{\text{st}}^V \mathbf{skip}$  in case  $v_{\text{active}} \neq B$ .  $\square$

Define  $\pi(\varepsilon) = \varepsilon$  and  $\pi(s; ss) = \pi(s); \pi(ss)$ . We have:

**Lemma B.7.** *Let  $P$  be a sequential program with variables  $V$ . For every sequence of statements  $ss$  of  $P$  in a block  $B$*

1. if  $v_{\text{active}} = B$ , then  $\pi(ss) \sim_{\text{st}}^V ss$ ;
2. if  $v_{\text{active}} \neq B$ , then  $\pi(ss) \sim_{\text{st}}^V \mathbf{skip}$ .

*Proof.* By induction on the structure of  $ss$  using Proposition B.3 and employing (1) Lemma B.6(1) if  $v_{\text{active}} = B$  and (2) Lemma B.6(2) and Proposition B.4 if  $v_{\text{active}} \neq B$ .  $\square$

**Stutter Equivalence of Blocks.** To achieve full generality with regard to stutter equivalence, we first drop the requirement that if a block ends with **goto**  $B_1, \dots, B_n$ , then at most one of  $B_1, \dots, B_n$  is a loop head. Predication of blocks is now defined as in Table 4. Where in the second row,  $B$  is assumed to be the last node, according to the sort order, of loops  $L_1, \dots, L_n$  and where  $B_{\text{head}}^{L_1}, \dots, B_{\text{head}}^{L_n}$  are the heads of the loops. The identifiers  $B_{\text{back}}^{L_1}, \dots, B_{\text{back}}^{L_n}$  and  $B_{\text{exit}}$  are fresh in the second row of the table. Note that the head of a loop  $L$  is the block that is always sorted first among the blocks of  $L$  (as it dominates all other blocks in  $L$ ).

We have the following:

**Lemma B.8.** *Let  $P$  be a sequential program with variables  $V$ . For every block  $B$  of  $P$*

1. if  $v_{\text{active}} = B$ , then  $\pi(B) \sim_{\text{st}}^V B$ ;
2. if  $v_{\text{active}} \neq B$ , then  $\pi(B) \sim_{\text{st}}^V B_{\text{skip}}$ ,

where  $B_{\text{skip}}$  denotes any block with **skip** as its only statement. Moreover,

1. if  $v_{\text{active}} = B$  and  $B$  ends with **goto**  $B_1, \dots, B_n$ , then for every terminating trace of the statements in  $\pi(B)$  it holds that  $v_{\text{active}} \in \{B_1, \dots, B_n\}$  in the final store of the trace;
2. if  $v_{\text{active}} \neq B$ , then the value of  $v_{\text{active}}$  is the same for every store in every trace of the statements in  $\pi(B)$ .

*Proof.* Let  $\sigma$  be a store of  $P$  and observe that  $\sigma|_V = \sigma$ . Moreover, observe that the only trace of **skip** is the terminating trace  $\sigma$ .

Assume  $ss$  is defined as

$$\begin{aligned} v_{\text{next}} &:= \{B_1, \dots, B_n\}; \\ v_{\text{active}} &:= (v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}}; \end{aligned}$$

We consider  $v_{\text{active}} = B$  and  $v_{\text{active}} \neq B$  in turn, where we assume  $\hat{\sigma}$  is a store of  $\pi(P)$  with  $\hat{\sigma}|_V = \sigma$ .

Original form	Predicated form
$B : ss$ $\mathbf{goto} B_1, \dots, B_n;$ <i>(B is not the last node of any loop according to the sort order)</i>	$B : \pi(ss)$ $v_{\text{next}} := \in \{B_1, \dots, B_n\};$ $v_{\text{active}} := (v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}};$ $\mathbf{goto next}(B);$
$B : ss$ $\mathbf{goto} B_1, \dots, B_n;$ <i>(B is the last node of some loops according to the sort order)</i>	$B : \pi(ss)$ $v_{\text{next}} := \in \{B_1, \dots, B_n\};$ $v_{\text{active}} := (v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}};$ $\mathbf{goto} B_{\text{back}}^{L_1}, \dots, B_{\text{back}}^{L_n}, B_{\text{exit}};$ $B_{\text{back}}^{L_1} : \mathbf{assume} v_{\text{active}} = B_{\text{head}}^{L_1};$ $\mathbf{goto} B_{\text{head}}^{L_1};$ $\vdots$ $B_{\text{back}}^{L_n} : \mathbf{assume} v_{\text{active}} = B_{\text{head}}^{L_n};$ $\mathbf{goto} B_{\text{head}}^{L_n};$ $B_{\text{exit}} : \mathbf{assume} \bigwedge_{i=1}^n (v_{\text{active}} \neq B_{\text{head}}^{L_i});$ $\mathbf{goto next}(B);$

Table 4: Predication of blocks

1. Let  $v_{\text{active}} = B$ . The only maximal traces of  $ss$  not terminating as *infeasible* are the traces  $\hat{\sigma} \hat{\tau}_{B_i} \hat{\upsilon}_{B_i}$  with  $0 \leq i \leq n$  where  $\hat{\tau}_{B_i} = \hat{\sigma}[v_{\text{next}} \mapsto B_i]$  and  $\hat{\upsilon}_{B_i} = \hat{\tau}_{B_i}[v_{\text{active}} \mapsto B_i]$ . As  $v_{\text{active}}$  does not occur in  $P$  and as  $v_{\text{next}}$  is fresh,  $\hat{\sigma} \upharpoonright_V = \hat{\tau}_{B_i} \upharpoonright_V = \hat{\upsilon}_{B_i} \upharpoonright_V$ . Hence, as  $\sigma$  is arbitrary,  $ss \sim_{\text{st}}^V \mathbf{skip}$  and stutter equivalence follows by Lemma B.7(1) and Propositions B.3 and B.4. The complete result follows once we observe that  $\hat{\upsilon}_{B_i}(v_{\text{active}}) = B_i$  for all  $0 \leq i \leq n$ .
2. Let  $v_{\text{active}} \neq B$ . The only maximal traces of  $ss$  not terminating as *infeasible* are the traces  $\hat{\sigma} \hat{\tau}_{B_i} \hat{\upsilon}_{B_i}$  with  $0 \leq i \leq n$  where  $\hat{\tau}_{B_i} = \hat{\sigma}[v_{\text{next}} \mapsto B_i]$  and  $\hat{\upsilon}_{B_i} = \hat{\tau}_{B_i}[v_{\text{active}} \mapsto \hat{\tau}_{B_i}(v_{\text{active}})] = \hat{\tau}_{B_i}$ . As  $v_{\text{next}}$  is fresh,  $\hat{\sigma} \upharpoonright_V = \hat{\tau}_{B_i} \upharpoonright_V = \hat{\upsilon}_{B_i} \upharpoonright_V$ . Hence, as  $\sigma$  is arbitrary,  $ss \sim_{\text{st}}^V \mathbf{skip}$  and stutter equivalence follows Lemma B.7(2) and Propositions B.3 and B.4. The complete result follows once we observe that the only assignment to  $v_{\text{active}}$  in  $\pi(B)$  occurs in the last statement of  $ss$  and that  $\hat{\upsilon}_{B_i}(v_{\text{active}}) = \hat{\sigma}(v_{\text{active}})$ .  $\square$

We also have the following:

**Lemma B.9.** *Let  $P$  be a sequential program with variables  $V$ . For every block  $B$  of  $P$  that occurs last in the sort order for loops  $L_1, \dots, L_n$*

1. *if  $v_{\text{active}} = B_{\text{head}}^{L_i}$  for some  $0 \leq i \leq n$ , then  $B_{\text{back}}^{L_i} \sim_{\text{st}}^V B_{\text{skip}}$  and every trace of  $B_{\text{exit}}$  is infeasible;*
2. *if  $v_{\text{active}} \neq B_{\text{head}}^{L_i}$  for all  $0 \leq i \leq n$ , then  $B_{\text{exit}} \sim_{\text{st}}^V B_{\text{skip}}$  and every trace of  $B_{\text{back}}^{L_i}$  is infeasible,*

where  $B_{\text{skip}}$  denotes any block with **skip** as its only statement. Moreover, in neither case is the value of  $v_{\text{active}}$  changed.

*Proof.* Let  $\sigma$  be a store of  $P$  and observe that  $\sigma \upharpoonright_V = \sigma$ . Moreover, observe that the only trace of **skip** is the terminating trace  $\sigma$ . We consider the two cases in turn, where  $\hat{\sigma}$  is a store of  $\pi(P)$  such that  $\hat{\sigma} \upharpoonright_V = \sigma$  and

$$e \doteq \bigwedge_{i=1}^n (v_{\text{active}} \neq B_{\text{head}}^{L_i}).$$

1. If  $v_{\text{active}} = B_{\text{head}}^{L_i}$  for some  $0 \leq i \leq n$ , then  $\hat{\sigma}(v_{\text{active}} = B_{\text{head}}^{L_i})$  holds and the only trace of **assume**  $v_{\text{active}} = B_{\text{head}}^{L_i}$  is the terminating trace  $\hat{\sigma}$ . The only trace of **assume**  $e$  is *infeasible*. As  $\hat{\sigma} \upharpoonright_V = \sigma$ , the result follows.
2. If  $v_{\text{active}} \neq B_{\text{head}}^{L_i}$  for all  $0 \leq i \leq n$ , then  $\hat{\sigma}(e)$  holds and the only trace of **assume**  $e$  is the terminating trace  $\hat{\sigma}$ . For all  $1 \leq i \leq n$ , the only trace of **assume**  $v_{\text{active}} = B_{\text{head}}^{L_i}$  is *infeasible*. As  $\hat{\sigma} \upharpoonright_V = \sigma$ , the result follows.  $\square$

**Stutter Equivalence of Sequential Programs.** Predication of a sequential program  $P$  proceeds by (1) sorting all blocks as described earlier, (2) predicating each block, and (3) inserting the assignment  $v_{\text{active}} := \text{Start}$  at the beginning of  $\pi(\text{Start})$ .

The addition of  $v_{\text{active}} := \text{Start}$  ensures that the statements from  $\pi(\text{Start})$  are always reduced first. Observe that if  $\pi(\text{Start})$  is part of a loop, it can only be the head and, hence, by the construction in the second row of Table 4, it is guaranteed that  $v_{\text{active}} = \text{Start}$  in case execution returns to  $\pi(\text{Start})$ . Call  $\pi(\text{Start})$  with  $v_{\text{active}} := \text{Start}$  at the beginning simply  $\pi(\text{Start})$  from here on. We have:

**Lemma B.10.** *Let  $P$  be a sequential program with variables  $V$ . Then,  $\pi(\text{Start}) \sim_{\text{st}}^V \text{Start}$ . Moreover, if  $\text{Start}$  ends with **goto**  $B_1, \dots, B_n$ , then for every terminating trace of the statements of  $\pi(\text{Start})$  it holds that  $v_{\text{active}} \in \{B_1, \dots, B_n\}$  in the final store of the trace.*

*Proof.* Let  $\sigma$  be a store of  $P$  and let  $\hat{\sigma}$  be a store of  $\pi(P)$  such that  $\hat{\sigma} \upharpoonright_V = \sigma$ . The only trace of  $v_{\text{active}} := \text{Start}$  is the terminating trace  $\hat{\sigma} \hat{\tau}$  with  $\hat{\tau} = \hat{\sigma}[v_{\text{active}} \mapsto \text{Start}]$ . Hence, as  $\hat{\sigma} \upharpoonright_V = \hat{\tau} \upharpoonright_V = \sigma$  and  $\sigma$  is arbitrary, we have  $(v_{\text{active}} := \text{Start}) \sim_{\text{st}}^V \text{skip}$ . Reduction of the remaining of  $\pi(\text{Start})$  starts with  $v_{\text{active}} = \text{Start}$  and is, thus, stutter equivalent to  $\text{Start}$  by Lemma B.8, where  $v_{\text{active}} \in \{B_1, \dots, B_n\}$  in the final store of each terminating trace. The complete result now follows by Proposition B.3.  $\square$

We can now state the main theorem of this section:

**Theorem B.11.** *If  $P$  is a sequential program with variables  $V$ , then  $\pi(P) \sim_{\text{st}}^V P$ .*

To prove the theorem, we need the following observation, which is the essence of Lemma 5.6:

**Lemma B.12.** *Let  $P$  be a sequential program with variables  $V$ . For each block  $B$ , if  $\sigma$  is a store of  $P$  and  $\hat{\sigma}$  is a store of  $\pi(P)$  such that  $\hat{\sigma} \upharpoonright_V = \sigma$  and  $\hat{\sigma}(v_{\text{active}}) = B$ , then*

1. *if the reduction of  $B$  is immediately followed by reduction of a block  $C$ , then there exists a reduction of  $\pi(B)$  such that  $v_{\text{active}} = C$  at the end of  $\pi(B)$  and eventually  $\pi(C)$  is reduced. Moreover, every other reduction of  $\pi(B)$  with  $v_{\text{active}} = C$  at the end of  $\pi(B)$  is *infeasible*;*

2. if the reduction of  $\pi(B)$  ends with  $v_{\text{active}} = C$ , then there exists a reduction of  $B$  that is immediately followed by the reduction of a block  $C$ .

*Proof.* We consider each of the cases in turn.

1. If the reduction of  $B$  is immediately followed by reduction of  $C$ , then  $B$  ends in **goto**  $\dots C \dots$ . Hence, by Lemma B.8 (and Lemma B.10 in case  $B = \text{Start}$ ), there are terminating traces of the statements of  $B$  and, hence, reductions of  $B$  such that  $v_{\text{active}} = C$  at the end of  $\pi(B)$ .  
Consider the the CFG of  $P$ , clearly there is an edge from  $B$  to  $C$ . If the edge is *not* a back edge, then the first requirement on the sort order of blocks in  $P$  ensures that  $B \leq C$ . By construction of the **gotos** in  $\pi(P)$  and the observations regarding  $v_{\text{active}}$  in Lemmas B.8 and B.9 (Lemmas B.10 and B.9 in the case of  $B = \text{Start}$ ), the result follows.  
If the edge from  $B$  to  $C$  is a back edge, say of a loop  $L_i$ , then it follows by the same reasoning as above that eventually the block of  $L_i$  that is sorted last is reduced. By definition of  $B_{\text{back}}^{L_i}$  and Lemma B.9, the result now follows.
2. Immediate by Lemma B.8 (Lemma B.10 in case  $B = \text{Start}$ ), reasoning by contradiction.  $\square$

We can now prove Theorem B.11:

*Proof.* By Lemma B.12 we have for every reduction of  $P$  reducing in order the blocks  $B_1, B_2, \dots$  that there exists a reduction of  $\pi(P)$  such that the value of  $v_{\text{active}}$  is in turn equal to  $B_1, B_2, \dots$ , and vice versa. Reasoning by contradiction, it follows by Proposition B.3 and Lemmas B.8, B.9, and B.10 that for every trace of  $P$  we can construct a trace of  $\pi(P)$  that is stutter equivalent, and vice versa.  $\square$

## B.2 Procedure Calls

Define the *main routine* of a sequential program as the the set of blocks that occur in the CFG that has *Start* as the root node. We extend our kernel programming language with procedures. Grammar-wise this means that *Programs* will be defined as follows [9]:

$$\begin{aligned} \text{Program} &::= \text{Procedure}^* \text{Block}^+ \\ \text{Procedure} &::= \mathbf{proc} \text{ProcId}(\mathbf{in} : \text{Var}^*, \mathbf{out} : \text{Var}^*) \text{BlockId}; \\ \text{Block} &::= \text{BlockId} : \text{Stmts} \text{Transfer}; \\ \text{Transfer} &::= \mathbf{return} \mid \mathbf{goto} \text{BlockId}^+ \end{aligned}$$

where we assume that the CFGs of all procedures and of main routine are disjoint. That is, for two distinct procedures defined by *BlockIds*  $B$  and  $C$ , we have that the blocks reachable from  $B$  are disjoint from the blocks reachable from  $C$  and that both are disjoint from the blocks reachable from *Start*.

For each procedure, **in** specifies a list of input parameters and **out** specifies a list of output parameters, both lists should be without duplicates. Both **return** and **goto** transfer control to some other block, where it is assumed that **return** does not occur in the main routine.

*Remark B.13.* We disallow **return** from occurring in the main routine, as this avoids having to deal with a host of special cases. Any program that does include a **return** in the main routine is easily rewritten to one without: Simply replace the **return** by a **goto** *End*.

Observe that, symmetrically to disallowing **return** in the main routine, a program cannot terminate from within a procedure: The *End* block is part of the main routine and, hence, cannot occur in the CFG of any procedure. Again, this avoids having to deal with a host of special cases. Any program which does allow termination from within a procedure can be rewritten into a program which does not exhibit this behaviour, albeit this is a harder than in the case of a **return** occurring in the main routine.

Statements are defined as before, except that we also have **call**-statements:

$$Var^* := \text{call } ProcId(Expr^*)$$

where the number of expressions is equal to the number of input parameters of the procedure and where the number of variables on the left of the assignment is equal to the number of output parameters. The variables are assigned values upon return from the procedure; no duplicates may occur among the variables and all variables must be thread-private.

**Semantics.** We describe the operational and denotational semantics of programs with procedures.

*Operational Semantics.* To give an operational meaning to procedure calls, we replace the private stores in the operational semantics from Figs. 4a and 4b by *stacks*. A stack is defined as a list of *stack frames*, where  $S \cdot F$  denotes a stack with the frame  $F$  on top. Denoting a finite sequence of variables  $v_1, \dots, v_n$  by  $v$ , a stack frame is a tuple  $(\sigma, v, v_{\text{out}}, b)$  with (a)  $\sigma$  a store that maps the *local variables* of a procedure to values, (b)  $v$  the names of variables that should be assigned values upon return from the procedure, (c)  $v_{\text{out}}$  the names of output parameters of the procedure, and (d)  $b$  a program fragment, which defines how reduction must continue upon return from the current procedure. Observe that a stack frame does not explicitly map input parameters to values. Instead, input parameters are considered to be local variables of a procedure. Output parameters are also local variables. Properly returning the values of output parameters is achieved by retaining in the stack frame the names of the parameters and the variables they need to be assigned to.

As before reduction of threads and sequential programs starts from *Start*, but the initial private stores are replaced by stacks containing a single stack frame  $(\sigma, \varepsilon, \varepsilon, \varepsilon)$  where  $\sigma$  is a store that maps to values the *global variables* of the thread/sequential program and where the first two occurrences of  $\varepsilon$  denote an empty sequence of variables and where the last occurrence of  $\varepsilon$  denotes an empty sequence of statements. The empty sequence of statements in the stack frame will never be reduced, as we assume that **return** does not occur in the main routine.

We denote the simultaneous update of multiple variables by  $\sigma[v \mapsto val]$ , where we assume that the variables in  $v$  are all distinct. Projecting a stack  $S_\sigma$  to the list of its stores

$$\begin{array}{c}
\frac{(\mathbf{proc} p(\mathbf{in} : \mathbf{v}_{\mathbf{in}}, \mathbf{out} : \mathbf{v}_{\mathbf{out}}) B_p) \in P \quad \tau_t \text{ a store} \quad F = (\mathbf{v}, \mathbf{v}_{\mathbf{out}}, b, \tau_t[\mathbf{v}_{\mathbf{in}} \mapsto \sigma(\mathbf{e})])}{P \vdash \langle \sigma, S_{\sigma_t}, \mathbf{v} := \mathbf{call} p(\mathbf{e}); b \rangle \xrightarrow{(\sigma, \sigma_t)} \langle \sigma, S_{\sigma_t} \cdot F, B_p \rangle} \text{CALL} \\
\frac{F = (\mathbf{v}, \mathbf{v}_{\mathbf{out}}, b, \tau_t) \quad \mathbf{val} = (\sigma, \sigma_t \cdot \tau_t)(\mathbf{v}_{\mathbf{out}})}{P \vdash \langle \sigma, S_{\sigma_t} \cdot F, \mathbf{return}; \rangle \xrightarrow{(\sigma, \sigma_t, \tau_t)} \langle \sigma, S_{\sigma_t}[\mathbf{v} \mapsto \mathbf{val}], b \rangle} \text{RETURN}
\end{array}$$

Fig. 8: Operational semantics of procedure calls and returns

$\sigma = \sigma_1 \cdot \dots \cdot \sigma_n$ , we denote by  $\sigma(\mathbf{e})$  evaluation of the expressions  $\mathbf{e}$  under  $\sigma$ , where the global variables are evaluated according to  $\sigma_1$  and the local variables according to  $\sigma_n$  (all other stores on the stack are inaccessible during evaluation). By  $S_\sigma[\mathbf{v} \mapsto \mathbf{val}]$  with  $\sigma = \sigma_1 \cdot \dots \cdot \sigma_n$  we denote the simultaneous update of the (thread-private) global variables in  $\sigma_1$  and the (thread-private) local variables in  $\sigma_n$  as specified by  $\mathbf{v}$ . In case of evaluation and updates, we assume that, if  $v$  is the name of *both* a local and a global variable, then the local variable  $v$  *hides* the global variable  $v$ .

The operational semantics of the kernel language extended with procedures is now defined as in Figs. 4a and 4b, where each store  $\sigma$  replaced by a stack and where each non-empty  $\rightarrow$ -label becomes the projection to the list of stores of the stack that occurs on the left of the arrow. The additional rules for **call** and **return** are given in Fig. 8. Observe that CALL is defined on the block level and not on the (possibly expected) statement level, because we need to take into account the continuation after the return from the procedure.

*Denotational Semantics.* The denotational semantics of kernels and sequential programs with procedures is defined as earlier, where thread-private stores are replaced by lists of stores. Observe that Propositions B.3 and B.4 still hold under introduction of procedures and when replacing stores by lists of stores.

Given a list of stores  $\sigma$  and a set of variables  $V$ , we define the *restriction* of  $\sigma$  to  $V$ , denoted  $\sigma|_V$ , by restricting each store in  $\sigma = \sigma_1 \cdot \dots \cdot \sigma_n$  to  $V$ , i.e.,  $\sigma|_V = \sigma_1|_V \cdot \dots \cdot \sigma_n|_V$ . Stutter equivalence is defined as before, *mutatis mutandis*.

**Predication of Procedures.** To predicate programs with procedures, we make the following two assumptions with regard to the programs we wish to predicate (to avoid a host of special cases):

1. each procedure call occurs in block of its own, i.e. the body of a block in which a call to a procedure  $p$  occurs always has the form  $\mathbf{v} := \mathbf{call} p(\mathbf{e}); t$ , where  $t$  only transfers control.
2. Each procedure has a unique block in which **return** occurs and the sequence of statements in the body of this block is *empty*.

A program is easily transformed into one which satisfies the above requirements.

Blocks are sorted *per* procedure following the earlier requirements. We require that the unique block with **return** is always sorted last to avoid premature return from

procedures. Moreover, we assume that blocks with **call**-statements are not sorted last among the blocks in a loop. If a block with a call does occur last, then we replace the **goto**  $B_1, \dots, B_n$  of the block by **goto**  $B_{\text{last}}$  with **goto**  $B_{\text{last}}$  fresh and we define

$$B_{\text{last}} : \mathbf{goto} B_1, \dots, B_n.$$

Observe that  $B_{\text{last}}$  will be part of the loop and sorted after the block with the **call**.

Predication replaces each procedure definition

$$\mathbf{proc} p(\mathbf{in} : v_{\text{in}}, \mathbf{out} : v_{\text{out}}) B_p$$

with

$$\mathbf{proc} p(\mathbf{in} : v_{\text{in}}, v_{\text{active}}, \mathbf{out} : v_{\text{out}}) B_p.$$

That is, we specify an input parameter for the initial value of  $v_{\text{active}}$  in  $p$ . Adding  $v_{\text{active}}$  as an input parameter avoids having to modify the initial block  $B_p$  of every procedure  $p$  in the same way we modified *Start*. Moreover, adding  $v_{\text{active}}$  as an input parameter will be of use in our treatment of our lock-step programs — it allows us to call a procedure as a no-op in case it does not need to be called.

Blocks with procedure calls are now predicated as in the top row of Table 5, where  $B_{\text{call}}$  and  $B_{\text{nocall}}$  are fresh. The block of a procedure in which the **return** occurs, which is unique and contains only a **return** by the restrictions from above, is simply copied during predication, as depicted in the bottom row of Table 5. Observe that reduction of the block is not restricted by the value of  $v_{\text{active}}$ . Hence, any sort order in which the **return** block does not occur last will result in premature return from the procedure.

Original form	Predicated form
$B : v := \mathbf{call} p(e);$ $\mathbf{goto} B_1, \dots, B_n;$	$B : \mathbf{goto} B_{\text{call}}, B_{\text{nocall}};$ $B_{\text{call}} : \mathbf{assume} v_{\text{active}} = B;$ $v := \mathbf{call} p(e, B_p);$ $v_{\text{next}} := \in \{B_1, \dots, B_n\};$ $v_{\text{active}} := (v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}};$ $\mathbf{goto} \mathit{next}(B);$ $B_{\text{nocall}} : \mathbf{assume} v_{\text{active}} \neq B;$ $\mathbf{goto} \mathit{next}(B);$
$B : \mathbf{return};$	$B : \mathbf{return};$

Table 5: Predication of procedure calls and returns

Predication of a sequential program  $P$  with procedures now proceeds in five steps: (1) sort the blocks of the main routine, (2) sort the blocks of each procedure, with the block from the procedure definition sorted first and with the unique block with the **return**-statement sorted last, (3) change the definition of each procedure, adding  $v_{\text{active}}$  as input parameter, (4) predicate each block as in Tables 4 and 5, with statements predicated as in Table 1, and (5) insert the assignment  $v_{\text{active}} := \mathit{Start}$  at the beginning of the predicated *Start* block.

**Stutter Equivalence of Sequential Programs with Procedures.** Ignoring again the shared store and observing that the results below go through if a shared store is present, mutatis mutandis, we next prove stutter equivalence for sequential programs with procedures.

Let the variables of a program be precisely those variables that either occur (a) in the main routine, (b) as input or output parameters of a procedure, or (c) in the blocks of procedures. Replaying the proofs of Lemmas B.6, B.7, and B.8, with the stores in the proofs replaced by stacks, we obtain:

**Lemma B.14.** *Let  $P$  be a sequential program with variables  $V$ . For every block  $B$  of  $P$  without any **call**-statement*

1. *if  $v_{\text{active}} = B$ , then  $\pi(B) \sim_{\text{st}}^V B$ ;*
2. *if  $v_{\text{active}} \neq B$ , then  $\pi(B) \sim_{\text{st}}^V B_{\text{skip}}$ ,*

where  $B_{\text{skip}}$  denotes any block consisting of a single **skip** statement. Moreover,

1. *if  $v_{\text{active}} = B$  and  $B$  ends with **goto**  $B_1, \dots, B_n$ , then for every terminating trace of the statements in  $\pi(B)$  it holds that  $v_{\text{active}} \in \{B_1, \dots, B_n\}$  in the final list of stores of the trace;*
2. *if  $v_{\text{active}} \neq B$ , then the value of  $v_{\text{active}}$  is the same for every list of stores in every trace of the statements in  $\pi(B)$ .*

In similar vein, we can replay Lemma B.9, where we should consider loops in both the main routine and in all procedures. We also have:

**Lemma B.15.** *Let  $P$  be a sequential program with variables  $V$ . For every block  $B$  of  $P$  with any **call**-statement such that  $B$  ends in **goto**  $B_1, \dots, B_n$  and where  $s$  is defined as **assume**  $v_{\text{active}} = B$ ; and where  $b$  is defined as*

$$\begin{aligned} v_{\text{next}} &:= \in \{B_1, \dots, B_n\}; \\ v_{\text{active}} &:= (v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}}; \\ \mathbf{goto} &\text{ next}(B); \end{aligned}$$

1. *if  $v_{\text{active}} = B$ , then  $s \sim_{\text{st}}^V \mathbf{skip}$  and  $b \sim_{\text{st}}^V B_{\text{skip}}$  and every trace of  $B_{\text{nocall}}$  is infeasible;*
2. *if  $v_{\text{active}} \neq B$ , then  $B_{\text{nocall}} \sim_{\text{st}}^V B_{\text{skip}}$  and every trace of  $s$  is infeasible,*

where  $B_{\text{skip}}$  denotes any block with **skip** as its only statement. Moreover,

1. *if  $v_{\text{active}} = B$ , then for every terminating trace of the statement in  $b$  it holds that  $v_{\text{active}} \in \{B_1, \dots, B_n\}$  in the final list of stores of the trace;*
2. *if  $v_{\text{active}} \neq B$ , then the value of  $v_{\text{active}}$  is the same for every list of stores in every trace of the statements in  $B_{\text{nocall}}$ .*

*Proof.* Let  $\sigma$  be a list of stores of  $P$  and observe that  $\sigma \upharpoonright_V = \sigma$ . Observe also that the only trace of **skip** is the terminating trace  $\sigma$ . We consider the two cases in turn, where we assume that  $\hat{\sigma}$  is a list of stores of  $\pi(P)$  with  $\hat{\sigma} \upharpoonright_V = \sigma$ :

1. If  $v_{\text{active}} = B$ , then  $\hat{\sigma}(v_{\text{active}} = B)$  holds. Hence, (a) the only trace of  $s$  is the terminating trace  $\hat{\sigma}$ , (b) the only maximal traces of the statement in  $b$  that are not *infeasible* are the terminating traces  $\hat{\sigma} \hat{\tau}_{B_i} \hat{v}_{B_i}$  with  $1 \leq i \leq n$  where  $\hat{\tau}_{B_i} = \hat{\sigma}[v_{\text{next}} \mapsto B_i]$  and  $\hat{v}_{B_i} = \hat{\tau}[v_{\text{active}} \mapsto B_i]$ , and (c) the only trace of **assume**  $v_{\text{active}} \neq B$  is *infeasible*. As  $v_{\text{next}}$  and  $v_{\text{active}}$  do not occur in  $P$ , we have  $\hat{\sigma}|_V = \hat{\tau}_{B_i}|_V = \hat{v}_{B_i}|_V$ . Thus, as  $\hat{\sigma}|_V = \sigma$ , we also have  $s \sim_{\text{st}}^V \text{skip}$  and  $b \sim_{\text{st}}^V B_{\text{skip}}$ . The result follows, once we observe that for all  $0 \leq i \leq n$  we have  $\hat{v}_{B_i}(v_{\text{active}}) = B_i$ .
2. If  $v_{\text{active}} \neq B$ , then  $\hat{\sigma}(v_{\text{active}} \neq B)$  holds. Hence, (a) the only trace of the statement **assume**  $v_{\text{active}} \neq B$  is the terminating trace  $\hat{\sigma}$  and (b) the only trace of  $s$  is *infeasible*. Thus, as  $\hat{\sigma}|_V = \sigma$ , we have  $B_{\text{nocall}} \sim_{\text{st}}^V B_{\text{skip}}$ . The result follows, once we observe that  $v_{\text{active}}$  is not assigned to in  $B_{\text{nocall}}$ .  $\square$

Before we arrive at the main theorem of this section, observe that we can replay Lemma B.10 in the current setting; by our earlier assumptions, no **call**-statement occurs in *Start*. Moreover, we have the following analogue of Lemma B.12, where  $\pi$  is inductively extended to lists of block bodies and where  $S_{\sigma, \gamma}$  denotes a stack which when projected to its stores and continuations yields, resp.,  $\sigma$  and  $\gamma$ .

**Lemma B.16.** *Let  $P$  be a sequential program with variables  $V$ . For every block  $B$ , if  $S_{\sigma, \gamma}$  is a stack of  $P$  and  $S_{\hat{\sigma}, \hat{\gamma}}$  is a stack of  $\pi(P)$  such that  $\hat{\sigma}|_V = \sigma$ ,  $\hat{\gamma} = \pi(\gamma)$ , and  $\hat{\sigma}(v_{\text{active}}) = B$ , then*

1. *if the reduction of  $B$  is immediately followed by reduction of a block  $C$ , then there exists a reduction of  $\pi(B)$  such that  $v_{\text{active}} = C$  at the end of  $\pi(B)$  and eventually  $\pi(C)$  is reduced. Moreover, every other reduction of  $\pi(B)$  with  $v_{\text{active}} = C$  at the end of  $\pi(B)$  is *infeasible*;*
2. *if the reduction of  $\pi(B)$  ends with  $v_{\text{active}} = C$ , then there exists a reduction of  $B$  that is immediately followed by the reduction of a block  $C$ .*

*Proof.* We consider each of the cases in turn.

1. Let the reduction of  $B$  immediately be followed by reduction of  $C$ . There are three cases to consider depending on the way in which control is transferred:
  - Let control be transferred by a **goto**. In this case, we can replay the proof of Lemma B.12 per procedure, *mutatis mutandis*.
  - Let control be transferred by a **call** to a procedure  $p$  which first reduces a block  $C$ . In this case,  $\pi(B)$  is either followed by reduction of  $B_{\text{call}}$  or  $B_{\text{nocall}}$ . If  $B_{\text{call}}$  is reduced, then the result follows by Lemma B.15 and the observation that  $\pi(p)$  is called with  $v_{\text{active}}$  set to  $C$ . If  $B_{\text{nocall}}$  is reduced, then reduction ends in *infeasible*, again by Lemma B.15.
  - Let control be transferred by a **return**. In this case, we have by definition of predication that the reduction of  $B$  is followed by a continuation of the form **goto**  $B_1, \dots, B_n$ . Hence, by Lemma B.15 and the assumption that  $\hat{\gamma} = \pi(\gamma)$ , we have that  $v_{\text{active}} = C$  at the end of the continuation in predicated form. Consider the CFG of the caller of the procedure  $p$  we are returning from, clearly there must be an edge to  $C$  from the block, say  $B'$ , with the call to  $p$ . Hence, the sort order on blocks in procedures ensures that  $B' \leq C$ . By construction

- of the **gotos**, the assumption that no back edges have  $B'$  as source, and the observation regarding  $v_{\text{active}}$  in Lemma B.14, the result now follows.
2. Let the reduction of  $\pi(B)$  end with  $v_{\text{active}} = C$ . There are again three cases to consider:
    - Let control be transferred by a **goto**. In this case, we can replay the proof of Lemma B.12 per procedure, *mutatis mutandis*.
    - Let control be transferred by a **call** to a procedure  $\pi(p)$  with  $v_{\text{active}}$  set to  $C$ . In this case, it follows by definition of predication that  $p$  first reduces a block  $C$  and, hence, the result is immediate.
    - Let control be transferred by a **return**. In this case, we have by definition of predication that reduction of  $B$  is followed by a continuation of the form **goto**  $B_1, \dots, B_n$ . The result now follows by Lemma B.15 and the assumption that  $\hat{\gamma} = \pi(\hat{\gamma})$ .  $\square$

The main theorem of this section now follows by replaying the proof of Theorem B.11, *mutatis mutandis*, where we employ Lemma B.16 instead of Lemma B.12.

**Theorem B.17.** *If  $P$  is a sequential program with variables  $V$ , then  $\pi(P) \sim_{\text{st}}^V P$ .*

### B.3 Stutter Equivalence of Interleaving Kernels

To ensure the correct operation of procedures in combination with barriers, we introduce an additional set of barrier variables:

- Every thread has a local variable  $v_{\text{call}}$  in each procedure (and the main routine). Each variable  $v_{\text{call}}$  is initialised as  $(-, -)$ . When a procedure  $p$  is called from a block  $B$ ,  $v_{\text{call}}$  is set to  $(p, B)$ , and it is reset to  $(-, -)$  upon return from  $p$ .

Moreover, every barrier variable introduced previously is assumed to be local. The additional constraint ensures that different threads follow the same sequence of calls to reach a barrier. Introducing these variables avoids having to inline all procedures, as is usually done by compilers of GPU kernels. Moreover, it allows us to handle recursion properly, which is not easily done with inlining.

In addition to the above we replace the  $\text{BARRIER}_F$  rule to properly handle non-terminating kernels:

$$\frac{\exists t : T_{\vec{\sigma}}|_t = \langle \beta_t, \sigma_t, \mathbf{barrier} \ e_t; b_t \rangle \wedge (\sigma, \sigma_t)(e_t) \quad \neg \forall \text{ maximal } (P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \rightarrow^* s) : s = \perp}{\neg \exists P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \rightarrow^* \langle \sigma, T_{\vec{\sigma}} \rangle : \forall t : T_{\vec{\sigma}}|_t = \langle \beta_t, \sigma_t, \mathbf{barrier} \ e_t; b_t \rangle \wedge (\sigma, \sigma_t)(e_t)} \text{BARRIER}_F$$

$$P \vdash \langle \sigma, T_{\vec{\sigma}} \rangle \xrightarrow{(\sigma, \vec{\sigma})} \mathcal{E}$$

Hence, a kernel terminate with *error* in case (a) one thread has reached a barrier, (b) not all traces end in *infeasible* irrespective of the barrier behaviour, and (c) it is impossible to reach a barrier where all barrier variables agree. In the rule,  $\rightarrow^*$  denotes the transitive, reflexive closure of  $\rightarrow$ . Note that the premise of  $\text{BARRIER}_F$  codifies an instance of the halting problem and, hence, is in general undecidable. However, in case we have a kernel that (a) is terminating, (b) has at its last statement a barrier, and (c) is well-formed, it follows easily that the above rule can be replaced by our original rule from Fig. 4d, albeit some extra reduction steps might be needed to have all threads reach a barrier.

**Stutter Equivalence of Interleaving Kernels.** To prove stutter equivalence of a kernel  $P$  with variables  $V$  (which are *not* barrier variables) and the predicated kernel  $\pi(P)$ , it is essential to observe that stutter equivalence of sequential programs, i.e. Theorem B.17, continues to hold under introduction of the rule

$$\frac{v \in V^* \quad val \in D^*}{P \vdash \langle S_\sigma, b \rangle \xrightarrow{\sigma} \langle S_\sigma[v \mapsto val], b \rangle} \text{UPDATE}$$

where the number of variables in  $v$  is equal to the number of values in  $val$  and where  $V$  consists solely of variables from  $P$  (no barrier variables or fresh variables from  $\pi(P)$  occur in  $V$ );

The above follows easily by replaying the proof of Theorem B.17. The UPDATE-steps in reductions are copied verbatim to the corresponding reductions of  $\pi(P)$ , and vice versa; no duplication or removal of such steps occurs. Observe that a reduction may now be infinite, because an infinite number of UPDATE-steps may occur at the end of the reduction. This should be taken into account when replaying the proof of Theorem B.17.

**Theorem B.18.** *If  $P$  is a kernel with variables  $V$  (not including barrier variables), then  $\pi(P) \sim_{\text{st}}^V P$ .*

*Proof.* Observe that adding a **skip**-statement before a **barrier**-statement in either a predicated and unpredicated kernel yields a kernel that is stutter equivalent to the original (predicated or unpredicated) kernel. Hence, it suffices to consider programs in which this additional **skip** is always present.

Let  $v_{\text{state}}$  be a fresh, shared variable and define  $\text{proj}_t$  for  $1 \leq t \leq TS$  to be a map from reductions in (interleaving) kernels to reductions in sequential programs:

- $\text{proj}_t$  maps  $\langle \sigma, T_{\vec{\sigma}} \rangle$  to  $\langle \sigma', S_{\sigma_t}, b'_t \rangle$  where  $T_{\vec{\sigma}}|_t = \langle S_{\sigma_t}, b_t \rangle$  and with
  1.  $\sigma'(v) = \sigma(v)$  for all  $v \in V$ ,
  2.  $\sigma'(v_{\text{state}}) = \langle \sigma, T_{\vec{\sigma}} \rangle[\bullet]_t$  for some constant  $\bullet$ ,
  3. and  $b'_t$  identical to  $b_t$  but with (a) each **skip** preceding a barrier replaced by an **assume** that states that the premise of one of the BARRIER-rules is satisfied and (b) each **barrier** replaced by an **assert** that states that the premise of either the BARRIER<sub>SKIP</sub>- or BARRIER<sub>S</sub>-rule is satisfied (all appropriately adapted to the introduction of  $v_{\text{state}}$ );
- $\text{proj}_t$  maps a step in a reduction to:
  1. an application of the ASSUME<sub>T</sub>-rule in case the THREAD<sub>B</sub>-rule was employed to reduce a **skip**-statement before a barrier and  $t$  was responsible for the step;
  2. an application of an ASSERT-rule in case a BARRIER<sub>B</sub>-rule was employed;
  3. the step in the premise of the applied rule in case either the THREAD<sub>B</sub>- or the THREAD<sub>E,I</sub>-rule was employed and  $t$  was responsible for the step;
  4. an application of the UPDATE-rule with  $v$  empty in case the THREAD<sub>T</sub>-rule was employed and  $t$  was responsible for the step;
  5. an application of the END-rule in case the TERMINATION-rule was employed;
  6. an appropriate application of the UPDATE-rule, otherwise;
 in each case,  $\sigma$  is replaced by  $\sigma'$  as defined above.

Let  $\Sigma \in \mathcal{D}(P)$  and  $\mathcal{D}(\rho) = \Sigma$ . If  $1 \leq t \leq TS$ , then, by the extension of Theorem B.17 with the SKIP-rule, we have  $\mathcal{D}(\text{proj}_t(\rho)) \sim_{\text{st}}^{V \cup \{v_{\text{state}}\}} \mathcal{D}(\rho')$  for some reduction  $\rho'$  of the *sequential* program  $\pi(P)$ , where barrier variables are not included in  $V$  and where the **skip ; barrier** is replaced as explained above for  $\text{proj}_t$  (all other statements are predicated as usual).

As UPDATE-steps are copied verbatim from  $\text{proj}_t(\rho)$  to  $\rho'$ , the inverse of  $\text{proj}_t$  exists for  $\rho'$  and, hence, there exists a trace  $\Sigma'$  that is stutter equivalent to  $\Sigma$  for  $V$  and that corresponds to a reduction that reduces  $\pi(\text{Start})$  for  $t$  and that reduces  $\text{Start}$  for all  $t' \neq t$ . By induction on the number of threads, applying the previous construction to each thread in turn, moving an increasing number of threads over to reduction of  $\pi(\text{Start})$ , it follows for every  $\Sigma \in \mathcal{D}(P)$  that there exists a  $\Sigma' \in \mathcal{D}(\pi(P))$  such that  $\Sigma \sim_{\text{st}}^V \Sigma'$ .

By a symmetric argument, there exists a  $\Sigma' \in \mathcal{D}(P)$  for every  $\Sigma \in \mathcal{D}(\pi(P))$  such that  $\Sigma \sim_{\text{st}}^V \Sigma'$ . Hence,  $\pi(P) \sim_{\text{st}}^V P$ .  $\square$

We have the following:

**Theorem B.19.** *Let  $P$  be kernel. A data race occurs in  $P$  iff a data race occurs in  $\pi(P)$  where during reduction of neither of the two statements causing the data race it holds that  $v_{\text{active}} \neq B$  where  $B$  is the block in which the statement occurs.*

*Proof.* To start, observe that every shared variable occurring in  $P$  also occurs in  $\pi(P)$ , and vice versa (all variables freshly introduced during the predication are thread-private). Hence, any data race that may occur in  $P$  must involve variables that also occur in  $\pi(P)$ , and vice versa.

Suppose that  $P$  has a data race. Then, by definition, a reduction  $\rho$  of  $P$  exists with a data race. Employ the construction from the proof of Theorem B.18 to obtain a reduction  $\rho'$  of  $\pi(P)$  with  $\rho' \sim_{\text{st}}^V \rho$  and  $V$  the non-barrier variables from  $P$ . By construction, we have for each block  $B$  that the variables accessed during reduction of  $B$  and  $\pi(B)$  are identical in case  $v_{\text{active}} = B$ . Moreover, no additional BARRIER<sub>s</sub> is inserted in between accesses. Hence,  $\rho'$  exhibits the same conflict as  $\rho$ .

Suppose that  $\pi(P)$  has a data race and consider the predicated form of the statements in Table 1. In case  $v_{\text{active}} \neq B$ , observe for call-statements that the only reductions of  $B_{\text{call}}$ -blocks immediately terminate as *infeasible*. For any other block freshly introduced during the predication, we have that only thread-private variables occur and, hence, reduction of these blocks cannot introduce any data races. By a similar argument as for data races in  $P$ , it now follows that for every reduction  $\rho$  in  $\pi(P)$  exhibiting some conflict, there exists a reduction  $\rho'$  in  $P$  that exhibits the same conflict.  $\square$

We have thus proved Theorem 5.5, as this combines Theorems B.18 and B.19.

#### B.4 Lock-Step Execution

In this section, we prove our soundness and completeness result. Before we can do so, we need to extend the lock-step construction for blocks, as we extended the handling of loops above. Moreover, we need to explain how we handle procedures in the lock-step construction.

**Lock-Step Construction for Blocks.** The lock-step construction for blocks is presented in Table 6a, where  $\phi(ss)$  denotes the lock-step construction for a sequence of statements  $ss$ , which is defined as  $\phi(\varepsilon) = \varepsilon$  and  $\phi(s; ss) = \phi(s); \phi(ss)$ , with  $\phi(s)$  the lock-step construction for a statement  $s$ .

In case a block *is not* sorted last among the blocks in a loop (see the top row of Table 6a) we simply duplicate updating  $v_{\text{active}}$ .

In case a block *is* sorted last among blocks in loops  $L_1, \dots, L_k, \dots, L_n$ , where we may assume without loss of generality that  $L_1 \subsetneq \dots \subsetneq L_k \subsetneq \dots \subsetneq L_n$ , we need to check whether:

1. we need to jump back to the head of a loop  $L_i$  for one of the threads (otherwise that thread will exit the loop prematurely), and whether
2. if we do need to jump back, that we do not need to jump back to the head of a loop  $L_j$  with  $1 \leq j < i$  for any thread (which would be a loop nested inside  $L_i$ ).

In addition, we only continue to reduce the next block in the sort order in case all loops have been exited by all threads. The case is presented in the bottom row of Table 6a.

**Lock-Step Construction for Procedures.** For lock-step execution of procedures, we need to ensure that if one thread calls a procedure then all threads do so, even if not all threads would have normally called the procedure. Of course, if a thread would not have normally called a procedure, we must enforce that the behaviour of the thread in the lock-step program is stutter equivalent to skipping the procedure.

Assuming that the block containing the unique **return** from a procedure  $p$  is  $End_p$  (cf. Sect. B.2), the above can be achieved by changing the  $B_{\text{call}}$  block from Table 5 to:

$$\begin{aligned} B_{\text{call}} : & \mathbf{v}_{\text{out}} := \mathbf{call} p(e, (v_{\text{active}} = B) ? B_p : End_p); \\ & \mathbf{v} := (v_{\text{active}} = B) ? \mathbf{v}_{\text{out}} : \mathbf{v}; \\ & v_{\text{next}} := \in \{B_1, \dots, B_n\}; \\ & v_{\text{active}} := (v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}}; \\ & \mathbf{goto} C; \end{aligned}$$

where  $v_{\text{out}}$  is a fresh variable. Observe that all statements in the called procedure  $p$  will effectively become no-ops, as the local  $v_{\text{active}}$  of the procedure is set to  $End_p$  and as the  $End_p$  block only returns.

Unfortunately, the above proposal for  $B_{\text{call}}$  is not stutter equivalent to the version from Table 5: If a recursive call to a procedure  $p$  occurs, then infinite recursion now occurs (calling  $p$  with  $End_p$  each time), while this might not have been the case in the original program. To avoid this problem, we ensure below that a procedure is only called in case it must be called by *at least* one thread.

The lock-step construction replaces each definition

$$\mathbf{proc} p(\mathbf{in} : v_{\text{in}}, v_{\text{active}}, \mathbf{out} : v_{\text{out}}) B_p,$$

with

$$\mathbf{proc} p(\mathbf{in} : \langle v_{\text{in}, t}, v_{\text{active}, t} \rangle_{t=1}^{TS}, \mathbf{out} : \langle v_{\text{out}, t} \rangle_{t=1}^{TS}) B_p.$$

Thus, similarly to the lock-step construction for assignment statements, we introduce a copy  $v_t$  of each parameter  $v$  for each thread  $t$ .

Predicated form	Lock-step form
$B$ : $ss$ $v_{\text{next}} \in \{B_1, \dots, B_n\}$ ; $v_{\text{active}} := (v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}}$ ; <b>goto</b> $next(B)$ ;	$B$ : $\phi(ss)$ $\langle v_{\text{next},t} \rangle_{t=1}^{TS} \in \{B_1, \dots, B_n\}^{TS}$ ; $\langle v_{\text{active},t} \rangle_{t=1}^{TS} := \langle (v_{\text{active},t} = B) ? v_{\text{next},t} : v_{\text{active},t} \rangle_{t=1}^{TS}$ ; <b>goto</b> $next(B)$ ;
$B$ : $ss$ $v_{\text{next}} \in \{B_1, \dots, B_n\}$ ; $v_{\text{active}} := (v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}}$ ; <b>goto</b> $B_{\text{back}}^{L_1}, \dots, B_{\text{back}}^{L_n}, B_{\text{exit}}$ ; $B_{\text{back}}^{L_1}$ : <b>assume</b> $v_{\text{active}} = B_{\text{head}}^{L_1}$ ; <b>goto</b> $B_{\text{head}}^{L_1}$ ; $\vdots$ $B_{\text{back}}^{L_k}$ : <b>assume</b> $v_{\text{active}} = B_{\text{head}}^{L_k}$ ; <b>goto</b> $B_{\text{head}}^{L_2}$ ; $\vdots$ $B_{\text{back}}^{L_n}$ : <b>assume</b> $v_{\text{active}} = B_{\text{head}}^{L_n}$ ; <b>goto</b> $B_{\text{head}}^{L_n}$ ; $B_{\text{exit}}$ : <b>assume</b> $\bigwedge_{i=1}^n (v_{\text{active}} \neq B_{\text{head}}^{L_i})$ ; <b>goto</b> $next(B)$ ;	$B$ : $\phi(ss)$ $\langle v_{\text{next},t} \rangle_{t=1}^{TS} \in \{B_1, \dots, B_n\}^{TS}$ ; $\langle v_{\text{active},t} \rangle_{t=1}^{TS} := \langle (v_{\text{active},t} = B) ? v_{\text{next},t} : v_{\text{active},t} \rangle_{t=1}^{TS}$ ; <b>goto</b> $B_{\text{back}}^{L_1}, \dots, B_{\text{back}}^{L_n}, B_{\text{exit}}$ ; $B_{\text{back}}^{L_1}$ : <b>assume</b> $\bigvee_{t=1}^{TS} (v_{\text{active},t} = B_{\text{head}}^{L_1})$ ; <b>goto</b> $B_{\text{head}}^{L_1}$ ; $\vdots$ $B_{\text{back}}^{L_k}$ : <b>assume</b> $(\bigvee_{t=1}^{TS} (v_{\text{active},t} = B_{\text{head}}^{L_k})) \wedge (\bigwedge_{i=1}^{k-1} v_{\text{active},t} \neq B_{\text{head}}^{L_i})$ ; <b>goto</b> $B_{\text{head}}^{L_2}$ ; $\vdots$ $B_{\text{back}}^{L_n}$ : <b>assume</b> $(\bigvee_{t=1}^{TS} (v_{\text{active},t} = B_{\text{head}}^{L_n})) \wedge (\bigwedge_{i=1}^{n-1} v_{\text{active},t} \neq B_{\text{head}}^{L_i})$ ; <b>goto</b> $B_{\text{head}}^{L_n}$ ; $B_{\text{exit}}$ : <b>assume</b> $\bigwedge_{i=1}^n (v_{\text{active},t} \neq B_{\text{head}}^{L_i})$ ; <b>goto</b> $next(B)$ ;

(a) Blocks

Predicated form	Lock-step form
$B$ : <b>goto</b> $B_{\text{call}}, B_{\text{nocall}}$ ; $B_{\text{call}}$ : <b>assume</b> $v_{\text{active}} = B$ ; $v := \text{call } p(e, B_p)$ ; $v_{\text{next}} \in \{B_1, \dots, B_n\}$ ; $v_{\text{active}} := (v_{\text{active}} = B) ? v_{\text{next}} : v_{\text{active}}$ ; <b>goto</b> $next(B)$ ; $B_{\text{nocall}}$ : <b>assume</b> $v_{\text{active}} \neq B$ ; <b>goto</b> $next(B)$ ; $B$ : <b>return ; </b>	$B$ : <b>goto</b> $B_{\text{call}}, B_{\text{nocall}}$ ; $B_{\text{call}}$ : <b>assume</b> $\bigvee_{t=1}^{TS} (v_{\text{active},t} = B)$ ; $\langle v_{\text{out},t} \rangle_{t=1}^{TS} := \text{call } p(\langle \phi(e) \rangle_t, (v_{\text{active},t} = B) ? B_p : \text{End}_p)^{TS}$ ; $\langle v_t \rangle_{t=1}^{TS} := \langle (v_{\text{active},t} = B) ? v_{\text{out},t} : v_t \rangle_{t=1}^{TS}$ ; $\langle v_{\text{next},t} \rangle_{t=1}^{TS} \in \{B_1, \dots, B_n\}^{TS}$ ; $\langle v_{\text{active},t} \rangle_{t=1}^{TS} := \langle (v_{\text{active},t} = B) ? v_{\text{next},t} : v_{\text{active},t} \rangle_{t=1}^{TS}$ ; <b>goto</b> $next(B)$ ; $B_{\text{nocall}}$ : <b>assume</b> $\bigwedge_{t=1}^{TS} (v_{\text{active},t} \neq B)$ ; <b>goto</b> $next(B)$ ; $B$ : <b>return</b> ;

(b) Procedures

Table 6: Lock-step construction for blocks and procedures

The lock-step construction for **call**-statements now combines the above proposal for  $B_{\text{call}}$  with a check that *at least* one thread needs to call the procedure. The construction is presented in the top row of Table 6b, where all thread-private variables are again duplicated for all threads (recall that all variables assigned to in the assignment of a call statement are thread-private).

As shown in the bottom row of Table 6b, a block with a **return** statement is left unchanged, as its only effect is to return from a procedure.

**Soundness and Completeness.** To prove soundness, we need to make one trivial change to our well-formed kernels. To see why a change is necessary, consider the following kernel:

```

Start : goto B1, B2;
B1   : assume true;
        barrier;
        goto End;
B2   : assume false;
        goto End;

```

If we apply our lock-step construction to this program, where we pick  $B_1 \leq B_2$ , then the condition in the **sync** that replaces the barrier will not evaluate to *true* in case we assign  $B_2$  to  $v_{\text{active},t}$  for some thread  $t$ . Although this assignment actually only yields infeasible traces in the interleaving semantics (due to the second condition in the premise of the  $\text{BARRIER}_F$ -rule), this behaviour is not captured by our lock-step construction, due to the sort order of blocks.

The above problem can be dealt with by “inserting relevant assumes” immediately before each goto. Suppose we have a block  $B$  in a *well-formed* kernel that ends in **goto**  $B_1, \dots, B_n$ , where  $e_1, \dots, e_n$  are the guards of the assumes of blocks  $B_1, \dots, B_n$  (a guard is equal to *true* in case no assume occurs). By predication, the **goto**-statement of  $B$  will be replaced by:

```

vnext := {B1, ..., Bn};
vactive := (vactive) ? vnext : vactive;
goto C;

```

Consider the following the replacement for the above code fragment:

```

vnext := {B1, ..., Bn};
assume (vactive = B) ⇒ ∧i=1n (vnext = Bi ⇒ ei);
vactive := (vactive) ? vnext : vactive;
goto C;

```

Here, we insert an early check for guards of the blocks whose reduction can follow reduction of  $B$ . Observe that this only reduces the length of infeasible traces, but does not affect the predicated kernel in any other way. Also observe that by well-formedness any reduction terminating as *infeasible* still can be transformed into one that does not terminate as *infeasible*.

The lock-step construction for the above assume is as for any other assume:

$$\mathbf{assume} \bigwedge_{t=1}^{TS} ((v_{\text{active},t} = B) \Rightarrow \bigwedge_{i=1}^n (v_{\text{next},t} = B_i \Rightarrow e_i));$$

With the above transformation, the observed problem with barriers can no longer occur, as the check for infeasibility is effectively moved in front of the barrier.

Assuming from here on that our lock-step programs include an **assume**-statement of the above kind between every assignment to  $v_{\text{next}}$  and  $v_{\text{active}}$ , we are now ready to prove soundness and completeness.

*Soundness.* We have the following with respect to data races:

**Theorem B.20.** *Let  $P$  be a well-formed kernel and let  $\phi(P)$  be the program obtained by applying the lock-step construction to  $P$ . The occurrence of a data race in  $\phi(P)$  implies the occurrence of a data race  $P$ .*

*Proof.* By Theorem B.19, it suffices to construct a reduction for  $\pi(P)$  — the predicated version of the kernel  $P$  — with an appropriate data race.

Let  $\rho$  be a reduction of  $\phi(P)$  and call a thread  $t$  *semi-active* in a step of  $\rho$ , if

1. whenever the step involves reduction of a statement from within a procedure  $p$ , then  $p$  was called with  $v_{\text{active},t} \neq \text{End}_p$ , and
2. whenever the step involves reduction a statement from a loop  $L_i$  or a block  $B_{\text{back}}^{L_i}$ , then  $v_{\text{active},t} \in L_i$ .

Call  $t$  *semi-inactive*, otherwise.

By an easy induction it follows that we can obtain a reduction of  $\pi(P)$  by replacing each reduction step in  $\rho$  by a sequence of steps which, instead of reducing one of the statements or transfers of control from the right-hand columns of Tables 3a, 6a, and 6b, reduces the corresponding statement or transfer of control from the left-hand column for *each* semi-active thread.

Care should be taken of the following:

- For an assignment to a shared variable, the application of the ASSIGN-step of a thread  $t$  should be last among all active threads in case the existential quantifier in the premise of  $\psi_T$  was instantiated with  $t$ .
- For a thread that should not call a procedure  $p$ , the transfer of control

**goto**  $B_{\text{call}}, B_{\text{nocall}}$

should reduce to  $B_{\text{nocall}}$  instead of  $B_{\text{call}}$ . In addition, instead of reducing the statement **assume**  $v_{\text{active}} = B$ , the statement **assume**  $v_{\text{active}} \neq B$  should be reduced.

- Reduction within  $B_{\text{call}}$ -blocks should be simplified in the obvious way.
- Each reduction of a **sync** statement should be appropriately replaced by either a **BARRIER<sub>SKIP</sub>**- or **BARRIER<sub>S</sub>**-step in case the expression of the **sync** evaluates to *true* and it should be replaced by a **BARRIER<sub>F</sub>**-step in case the expression evaluates to *false*. That the replacement is possible follows by following claim:

*Claim.* For the expression an **sync**-statement the following holds:

- if the expression evaluates to *true*, then the premise of either the **BARRIER<sub>SKIP</sub>**- or the **BARRIER<sub>S</sub>**-rule is satisfied, and

- if the expression evaluates to *false*, then the premise of  $\text{BARRIER}_F$  is satisfied.

*Proof.* We prove the two parts of the claim in turn.

- Suppose the expression evaluates to *true*. In this case, observe that a thread cannot progress from a semi-inactive state to a semi-active state while reduction is taking place inside a loop, as no statements from outside the loop occurs in between the statements from the loop (cf. the block sort order). The same holds while a thread is inside a procedure that was called in a semi-inactive state. Hence, in case the expression of a *sync*-statement evaluates to *true*, either no thread has reached a barrier, or it holds that the loop counters and  $v_{\text{call}}$  variables are identical across threads and, hence, that all threads have reached the same barrier by the lock-step execution. Whence, the premise of either the  $\text{BARRIER}_S$ - or the  $\text{BARRIER}_F$ -rule is satisfied.
- Suppose the expression evaluates to *false*. In this case, reason by contradiction and, thus, also suppose that for the reduction constructed up to the *sync*-statement either (a) every maximal extension ends in *infeasible*, or (b) can be extended to one that satisfies the premise of the  $\text{BARRIER}_S$ -rule. Consider a thread  $t$  with  $v_{\text{active},t} \neq B$  upon reaching the *sync*-statement. We consider the two cases in turn.
  - In case every maximal extension of the reduction is *infeasible*, the thread  $t$  must fail on the guard of the block that is chosen as alternative to the block with the barrier. However, due to the *assume*-statements introduced on page 38, it follows that the assumed interleaving reduction does not exist.
  - In case we can extend the reduction to one that satisfies the premise of the  $\text{BARRIER}_S$ -rule, the thread  $t$  can only reach the barrier if either the thread loops or calls an appropriate procedure later. Hence, the values of some loop variables or  $v_{\text{call}}$  variables will differ from those of the threads that already reached the barrier during lock-step execution. Whence, the premise of the  $\text{BARRIER}_S$ -rule can never be satisfied and we can apply the  $\text{BARRIER}_F$ -rule.  $\square$

It now follows immediately by the above construction and the definition of data races for kernels that each occurrence of a data race in  $\phi(P)$  implies the occurrence of a data race in  $\pi(P)$  such that during reduction of neither of the two statements causing the data race it holds that (a)  $v_{\text{active}} \neq B$  where  $B$  is the block in which the statement occurs and that (b) a  $\text{BARRIER}_S$ -step occurs in between the two statements.  $\square$

We immediately have the following corollary.

**Corollary B.21.** *Let  $P$  be a well-formed kernel and let  $\phi(P)$  be the program obtained by applying the the lock-step construction to  $P$ . If  $P$  is race free, then  $\phi(P)$  is race free.*

The reverse of the above corollary, i.e., completeness, does in general not hold. To see this, assume that each thread  $t$  has a private variable *tid* whose value is  $t$  and consider the following program, where  $v$  is a shared variable:

```

Start : goto B1, B2 ;
B1   : assume tid = 1 ;
        goto Start ;
B2   : assume tid ≠ 1 ;
        v := 1 ;
        goto End ;

```

Thread 1 loops forever, while every other thread assigns 1 to  $v$ . In case of the interleaving semantics, reductions exist in which thread 1 is ignored initially and in which every other thread is reduced until termination. Obviously, any such reduction has a data race in case there are at least three threads, as multiple threads assign 1 to the shared variable  $v$ . In the case of the lock-step semantics, every thread must first complete all loop iterations. As thread 1 never completes all loop iterations, the assignment to  $v$  is never reached and, hence, no data race occurs.

The above counterexample can be considered rather pathological due to its non-terminating nature (see also [11,12]). Hence, we next consider terminating programs.

*Soundness and Termination.* Our main soundness result is as follows:

**Theorem B.22 (Soundness).** *Let  $P$  be a well-formed kernel and let  $\phi(P)$  be the program obtained by applying the lock-step construction to  $P$ . If  $P$  is race free and terminating, then  $\phi(P)$  is race free and terminating. Moreover, if no data races occur, then for every terminating reduction of  $P$  there exists a terminating reduction of  $\phi(P)$  such that every shared variable  $v$  of  $P$  has the same value at the end of both reductions.*

To prove the theorem, we need the following observation:

**Proposition B.23.** *Let  $P$  be a kernel. If*

$$\langle \sigma, T_{\vec{\sigma}} \rangle \xrightarrow{(\sigma, \vec{\sigma})} \langle \tau, T_{\vec{\sigma}}[\langle S_{\tau}, c \rangle]_t \rangle \xrightarrow{(\tau, \vec{\tau})} \langle v, T_{\vec{\tau}}[\langle S_v, d \rangle]_{t'} \rangle$$

*is race free, with  $T_{\vec{\tau}} = T_{\vec{\sigma}}[\langle S_{\tau}, c \rangle]_t$  and  $1 \leq t \neq t' \leq TS$ , and the second step is either a  $\text{THREAD}_{\text{B-}}$ ,  $\text{THREAD}_{\text{T-}}$ , or  $\text{BARRIER}_{\text{SKIP}}$ -step, then there exists a race-free reduction*

$$\langle \sigma, T_{\vec{\sigma}} \rangle \xrightarrow{(\sigma, \vec{\sigma})} \langle \tau', T_{\vec{\sigma}}[\langle S_v, d \rangle]_{t'} \rangle \xrightarrow{(\tau', \vec{\tau}')} \langle v, T_{\vec{\chi}}[\langle S_{\tau}, c \rangle]_t \rangle$$

*with  $T_{\vec{\chi}} = T_{\vec{\sigma}}[\langle S_v, d \rangle]_{t'}$  and  $T_{\vec{\chi}}[\langle S_{\tau}, c \rangle]_t = T_{\vec{\tau}}[\langle S_v, d \rangle]_{t'}$ .*

*Proof.* Observe that only the  $\text{THREAD}_{\text{B-}}$ ,  $\text{THREAD}_{\text{T-}}$ , and  $\text{BARRIER}_{\text{SKIP}}$ -rules may have been employed as the first step in the two-step reduction. As the two steps do not race and as  $t \neq t'$ , the result is immediate by the fact that two threads cannot access each others' stacks. Observe by race-freeness that the premise of the rule applied in the second step of the original reduction are already satisfied in  $\langle \sigma, T_{\vec{\sigma}} \rangle$  and that the premise of the rule applied in the first step of the original reduction are still satisfied after performing  $\langle \sigma, T_{\vec{\sigma}} \rangle \rightarrow \langle \tau', T_{\vec{\sigma}}[\langle S_v, d \rangle]_{t'} \rangle$ .  $\square$

We can now prove Theorem B.22.

*Proof.* That  $\phi(P)$  is race free follows by Corollary B.21. That  $\phi(P)$  terminates follows by contradiction: If there exists an infinite and race free reduction for  $\phi(P)$ , then an infinite and race free reduction also exists for  $P$  by the construction in the proof of Theorem B.20.

For the second half of the theorem, let  $\rho$  be a terminating and race free reduction of  $P$ . Observe that before each  $\text{BARRIER}_S$ -step in  $\rho$  only  $\text{THREAD}_B$ - and  $\text{BARRIER}_{\text{SKIP}}$ -steps occur and that after the last  $\text{BARRIER}_S$ -step only  $\text{THREAD}_T$ -steps and a  $\text{TERMINATION}$ -step occur. By the above proposition we may re-order the steps of  $\rho$  before each  $\text{BARRIER}_S$ -step such that the eventual values of the shared variables are as in  $\rho$ . In fact, it is possible re-order the steps such that we obtain a reduction as the one constructed in the proof of Theorem B.20 — grouping reduction of the same statements and transfers of control from different threads. For suppose we cannot re-order in this way, then for one or more threads an insufficient number of steps occur before a certain barrier. However, in this case either (a) different threads have reached different barriers, (b) the number of loop iterations differs between threads, or (c) the sequence of procedure calls used to reach the barrier differs between threads. In each of these cases the premise of  $\text{BARRIER}_S$  is not satisfied, contradiction. The complete result now follows by induction applying the construction from Theorem B.20 in reverse. That the  $\text{BARRIER}_{\text{SKIP}}$ - and  $\text{BARRIER}_S$ -steps can be replaced by  $\text{SYNC}$ -steps follows by similar reasoning as in the proof of the claim on page 39.  $\square$

*Completeness.* We have the following:

**Theorem B.24 (Completeness).** *Let  $P$  be a well-formed kernel and let  $\phi(P)$  be the program obtained by applying the lock-step construction of  $P$ . If  $\phi(P)$  is race free and terminating, then  $P$  is race free and terminating. Moreover, if no data races occur, then for every terminating reduction of  $\phi(P)$  there exists a terminating reduction of  $P$  such that every shared variable  $v$  of  $P$  has the same value at the end of both reductions.*

To prove the theorem, we need the next observation, where we say that a data race with respect to a variable  $v$  is called:

- a *write-write conflict*, in case both threads update  $v$ ;
- a *read-write conflict*, in case one thread uses  $v$  to evaluate an expression and later along  $\rho$  another thread updates  $v$ ;
- a *write-read conflict*, in case one thread updates  $v$  and later along  $\rho$  another thread uses  $v$  to evaluate an expression.

**Proposition B.25.** *Let  $P$  be a well-formed kernel. If  $\rho \cdot \tau$  is a (partial) reduction of  $P$  where  $\tau$  consists of two steps and has a data race, then there exists a (partial) reduction  $\rho \cdot \tau'$  with  $\tau'$  consisting of at most two steps such that the first step of  $\tau'$  reduces the same statement in the same thread as the second step of  $\tau$  and either*

1.  $\tau'$  consists of a single step and terminates with error, or
2.  $\tau'$  has a data race and the second step of  $\tau'$  reduces the same statement in the same thread as the first step of  $\rho$ .

*Proof.* Observe that the two steps in  $\tau$  originate from different threads, otherwise no data race occurs in  $\tau$ .

Suppose the data race in  $\tau$  is a write-write conflict. In this case, the statements reduced in  $\tau$  are either assignments, havocs, returns from a procedures, or a combination of two thereof. As the steps originate from different threads, it is easily seen that the statements can be reduced in reverse order. No *error* can be introduced, as neither of the statements is an `assert` statement. A write-write conflict still occurs after reversal, as the updated variables do not change.

Suppose the data race in  $\tau$  is either a read-write or a write-read conflict. In this case, the write part of the conflict originates with either an assignment, a havoc, or a return from a procedure; the read part of the conflict originates with either an assignment, a call of a procedure, or an `assert` statement. In case the read part originates with an assignment or a call of a procedure, the reduction of the two statements can be reversed and a conflict still occurs, as no other variable is updated during the write part of the conflict. In case the read part originates with an `assert`, either the same reasoning as for the assignment applies, or the `assert` becomes the first statement that is being reduced and we terminate with *error*.

Note that `assume` statements do not need to be considered, as we have by well-formedness of  $P$  that no global variables occur in such statements.  $\square$

We can now prove Theorem B.24.

*Proof.* We reason by contradiction. Thus, suppose  $P$  is not terminating or race free. Then, either (a)  $P$  is race free and there exists a reduction  $\rho$  that is infinite or terminates with status *error*, or (b) there exists a reduction  $\rho$  of  $P$  with a data race.

In case of (a), employ Proposition B.23 to re-order the steps in  $\rho$  to obtain a reduction where the steps have the same order as those of the reduction constructed in the proof of Theorem B.20 — grouping reduction of the same statements and transfers of control from different threads. Before each `BARRIERS`-step, it suffices to re-order steps, by reasoning as in the proof of Theorem B.22. Following the last `BARRIERS`-step, steps may need to be added for some threads (or removed when dealing with the `BARRIERF`-rule) to ensure we conform to the grouping from the theorem (this may turn an infinite reduction into an *error* reduction or lead to an earlier *error*, but this is, of course, of no consequence). By assumption, the additional steps cannot introduce any data races. By applying the construction from Theorem B.20 in reverse (see also the proof of Theorem B.22), we obtain for  $\phi(P)$  an infinite reduction or an *error* reduction, contradiction.

In case of (b), employ Propositions B.23 and B.25 to re-order the steps of  $\rho$  to construct a (partial) reduction which either terminates with *error* or has a data race, and in which the steps have the same order as those of the reduction constructed in the proof of Theorem B.20 — grouping reduction of the same statements and transfers of control from different threads. Following the last `BARRIERS`-step, steps may need to be added for some threads (or removed when dealing with the `BARRIERF`-rule) to ensure we conform to the grouping from the theorem (this may turn an infinite reduction into an *error* reduction or lead to an earlier *error*, but this is, of course, of no consequence). By applying the construction from Theorem B.20 in reverse, we now obtain for  $\phi(P)$  a (partial) reduction which either terminates with *error* or has a data race. By well-formedness of  $P$ , this partial reduction can be extended to a maximal reduction of

$\phi(P)$  which is not *infeasible* and which either terminates with *error* or has a data race, contradiction.  $\square$

Combining Theorems B.22 and B.24 we obtain Theorem 5.2.

## C 2-Threaded Abstraction

Although considering lock-step programs instead of kernels offers the potential to significantly reduce the verification state space, due to the removal of interleavings, the space may still be too large to make verification feasible, due to the large number of threads which usually execute a GPU kernel. Observe, however, that a data race always involves an interaction between *two* threads, which leads to the idea of abstracting away all threads of a kernel except for two [7].

**Definition C.1 (2-Threaded Abstraction).** *Let  $P$  be a kernel. If  $1 \leq t_1 < t_2 \leq TS$ , then the 2-threaded abstraction of  $P$  with respect to  $t_1$  and  $t_2$ , denoted  $P_{t_1, t_2}$  is the kernel  $P$  where each barrier havocs every shared variable and with all reductions starting from*

$$\langle \sigma, \langle S_{\sigma_{t_1}}, Start \rangle, \langle S_{\sigma_{t_2}}, Start \rangle \rangle$$

instead of from

$$\langle \sigma, \langle S_{\sigma_1}, Start \rangle, \dots, \langle S_{\sigma_{TS}}, Start \rangle \rangle,$$

where  $S_{\sigma_t} = (\sigma_t, \varepsilon, \varepsilon, \varepsilon)$  with  $\sigma_t$  a store for  $1 \leq t \leq TS$ .

We have the following:

**Theorem C.2.** *Let  $P$  be a well-formed kernel. If for all  $1 \leq t_1 < t_2 \leq TS$  the program  $P_{t_1, t_2}$  is race free and terminating, then  $P$  is race free and terminating.*

*Proof.* Suppose  $P$  is either not terminating or not race free. Then, either (a)  $P$  is race free and there exists a reduction  $\rho$  that is infinite or terminates with *error*, or (b) there exists a reduction  $\rho$  of  $P$  with at least one data race.

In the case of (a), let  $t$  be a thread that is responsible for an infinite number steps of  $\rho$  or that is responsible for the termination with *error*. As  $P$  is assumed to be race free, we can remove from  $\rho$  all steps for which  $t$  is not responsible and that are not the responsibility of some arbitrary thread  $t' \neq t$ . The havocs performed at barriers ensure that the reduction  $\rho'$  so obtained is a reduction of  $P_{t, t'}$ , contradiction.

In the case of (b), consider the shortest (partial) reduction  $\rho'$  such that  $\rho'$  has a data race and is a prefix of  $\rho$ . Consider one of the data races in  $\rho'$  (there may be multiple, but all involve the last step of  $\rho'$ ) and suppose the two threads responsible for the data race are  $t$  and  $t'$ . Remove from  $\rho'$  all steps for which  $t$  and  $t'$  are not responsible. The havocking at barriers ensures that the reduction  $\rho'$  so obtained is a reduction of  $P_{t, t'}$ , which also has a data race (otherwise there was a shorter reduction with a data race) and which can be extended to a maximal reduction not terminating with *infeasible* by well-formedness of  $P$ , contradiction.  $\square$

Observe that the lock-step construction also applies to our abstracted programs  $P_{t_1, t_2}$ ; we simply need to replace every construct of the form  $C_{t=1}^{TS}$  with  $C_{t \in \{t_1, t_2\}}$ , where  $C$  is either  $\langle \rangle$ ,  $\wedge$ , or  $\vee$ . All facts established in the previous section continue to hold, mutatis mutandis. In particular, Theorem 5.2 continues to hold. Hence, we also have:

**Theorem C.3.** *Let  $P$  be a well-formed kernel. If for all  $1 \leq t_1 < t_2 \leq TS$  the lock-step program  $\phi(P_{t_1, t_2})$  is race free and terminating, then  $P$  is race free and terminating.*

Observe that the reverse of the above theorem does not hold due to the havocing introduced at each barrier.