Semantics of Remote Direct Memory Access: Operational and Declarative Models of RDMA on TSO Architectures

GUILLAUME AMBAL[*], Imperial College London, UK
BRIJESH DONGOL, University of Surrey , UK
HAGGAI ERAN, NVIDIA, Israel
VASILEIOS KLIMIS, Queen Mary University of London, UK
ORI LAHAV, Tel Aviv University, Israel
AZALEA RAAD, Imperial College London, UK

Remote direct memory access (RDMA) is a modern technology enabling networked machines to exchange information without involving the operating system of either side, and thus significantly speeding up data transfer in computer clusters. While RDMA is extensively used in practice and studied in various research papers, a formal underlying model specifying the allowed behaviours of concurrent RDMA programs running in modern multicore architectures is still missing. This paper aims to close this gap and provide semantic foundations of RDMA on x86-TSO machines. We propose three equivalent formal models, two operational models in different levels of abstraction and one declarative model, and prove that the three characterisations are equivalent. To gain confidence in the proposed semantics, the more concrete operational model has been reviewed by NVIDIA experts, a major vendor of RDMA systems, and we have empirically validated the declarative formalisation on various subtle litmus tests by extensive testing. We believe that this work is a necessary initial step for formally addressing RDMA-based systems by proposing language-level models, verifying their mapping to hardware, and developing reasoning techniques for concurrent RDMA programs.

CCS Concepts: • **Software and its engineering** → **Formal language definitions**; • **Theory of computation** → **Program semantics**; **Distributed computing models**; • **Hardware** → Testing with distributed and parallel systems.

Additional Key Words and Phrases: RDMA, Operational Semantics, Declarative Semantics, x86-TSO

ACM Reference Format:
Guillaume Ambal, Brijesh Dongol, Haggai Eran, Vasileios Klimis, Ori Lahav, and Azalea Raad. 2024. Semantics of Remote Direct Memory Access: Operational and Declarative Models of RDMA on TSO Architectures. *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 341 (October 2024), 68 pages. https://doi.org/10.1145/3689781

1 INTRODUCTION

Remote direct memory access (RDMA) technologies such as *InfiniBand* and *RDMA over Converged Ethernet* (RoCE) enable a machine to have *direct* read/write access to the memory of another machine over a network, bypassing the operating systems of both machines. This way, remote reads and writes are performed with far fewer CPU cycles, leading to high-throughput, low-latency

[*] Corresponding author

Authors' addresses: Guillaume Ambal, Imperial College London, UK, g.ambal@imperial.ac.uk; Brijesh Dongol, University of Surrey , UK, b.dongol@surrey.ac.uk; Haggai Eran, NVIDIA, Israel, haggaie@nvidia.com; Vasileios Klimis, Queen Mary University of London, UK, v.klimis@qmul.ac.uk; Ori Lahav, Tel Aviv University, Israel, orilahav@tau.ac.il; Azalea Raad, Imperial College London, UK, azalea.raad@imperial.ac.uk.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
© 2024 Copyright held by the owner/author(s).
ACM 2475-1421/2024/10-ART341
https://doi.org/10.1145/3689781

networking, which is especially useful in massively parallel computer clusters, e.g. for data centres, big data, and scientific computation. Thanks to implementations that offer higher performance at a comparable cost over traditional networking infrastructure (e.g. TCP/IP sockets) [Gerstenberger et al. 2018], RDMA has achieved widespread adoption as of 2018 [Shpiner et al. 2017] and has been rapidly adopted in modern data centres.

At the lowest level, RDMA networks directly interact with the hardware through calls to read (get) and write (put) operations to remote memory locations. As a result, programming RDMA systems is conceptually similar to shared memory systems of existing hardware architectures such as Intel-x86 or ARM. A key difference, however, is that when a machine encounters a remote operation, the CPU forwards it onto the *network interface card* (NIC), which subsequently handles the remote operation and its associated memory accesses without further CPU involvement.

There is a wide range of RDMA implementations, starting from the Virtual Interface Architecture (VIA) [Dunning et al. 1998], later adapted by InfiniBand [IBTA 2022] and RoCE [InfiniBand Trade Association (IBTA) 2018]. Other standards include iWarp [Recio et al. 2007] and Omni-Path [Birrittella et al. 2015]. Some transfer technologies without NICs, such as FireWire [Anderson 1999], can also be considered RDMA, but they do not scale to wide networks. In most implementations, an RDMA NIC implements the transport in hardware and is controlled by software through APIs such as Verbs [linux-rdma 2018] or libfabric [OpenFabrics 2016]. A NIC typically connects to the host CPU through an internal server fabric such as PCIe, though in some cases the NIC and compute cores can be more tightly integrated, e.g. in academic proposals [Novakovic et al. 2014] and in DPUs [NVIDIA Corporation 2021]. In addition, there have been proposals for next-gen fabrics to replace PCIe (e.g. CXL [Van Doren 2019]).

Our focus here is on the IB Verbs model defined by IBTA [2022], using PCIe as the internal fabric. It was designed for InfiniBand and reused for RoCE, the two most popular RDMA technologies. While the specification for the iWarp protocol is slightly more permissive, its main implementation (libfabric) follows the stronger semantics presented here. Lastly, the more recent Omni-Path fabric also has legacy support for IB Verbs.

The performance gains of RDMA, as well as its wide range of implementations, have led to a surge of both theoretical and practical RDMA research [Aguilera et al. 2019; Dan et al. 2016; Wei et al. 2015; Zhu et al. 2015]. However, as we discuss below, programming RDMA systems *correctly* is not straightforward, and the RDMA community would benefit greatly from formal models and rigorous techniques for reasoning about RDMA programs.

A key challenge lies in understanding the different degrees of concurrency and their interaction thereof. More concretely, a program may comprise threads that run over multiple nodes (machines) over the network (inter-node concurrency), with each node itself executing several threads (intra-node concurrency). As such, to understand the behaviour of a concurrent RDMA program, one must understand how remote and local operations on different nodes interact with one another. The problem is that local operations are handled by the CPU, while remote operations are handled by the NIC independently and in *parallel* to other CPU operations. Consequently, two sequential operations may not be executed in the intended (program) order, leading to surprising outcomes.

To understand the behaviour of RDMA programs, we must understand the order in which CPU and remote operations are executed, how they may be *reordered*, and how and when their effects are made visible to concurrent threads, be they on the same or different nodes. Specifically, much in the same way that the semantics of multi-processor hardware architectures such as Intel-x86, POWER, and ARM have been described via *formal consistency models* (a.k.a. *memory models*) [Alglave et al. 2021, 2014; Cho et al. 2021; Mador-Haim et al. 2012; Pulte et al. 2018, 2019; Raad et al. 2022; Raad and Vafeiadis 2018; Raad et al. 2020b, 2019b; Sarkar et al. 2011; Sewell et al. 2010], we should ideally describe the semantics of RDMA programs *formally*. Such formal models not only provide a rigorous underpinning for reasoning about the behaviour of programs, they have also been historically successful at identifying mistakes and ambiguities in the existing hardware reference manuals, as well as compilation bugs [Alglave et al. 2021, 2014; Lahav et al. 2017; Pulte et al. 2018; Raad et al. 2020b].

Unfortunately, the existing literature includes next to no work on the formal semantics of RDMA programs. Indeed, to our knowledge, the coreRMA model by Dan et al. [2016] is the only one to offer a formal description of RDMA programs. However, this work has four key shortcomings.

First, coreRMA assumes that CPU concurrency on each node is governed by the *sequential consistency* (SC) [Lamport 1979] model. This is an unrealistic assumption as no existing CPU architecture supports SC by default, and the two mainstream CPU architectures, Intel-x86 and ARM, are both subject to weaker models that exhibit behaviours not possible under SC. This is a significant gap since Intel-x86 and ARM architectures are ubiquitous.

Second, coreRMA describes the semantics *only declaratively* (i.e. via execution graphs that are subject to a set of axioms stipulating the absence of certain cycles in the graphs) and *not operationally* (via transitions that describe how the underlying hardware processes each operation and manipulates the memory). While declarative models are common in the literature of weak memory models and are more concise, operational models provide a more intuitive account of the hardware guarantees (as they prescribe a step-by-step mechanism for producing the behaviours of a program). Moreover, having two characterisations is useful not only for ensuring the accuracy of the formalism, but also because each formulation may be more useful for establishing different results. In particular, operational models are better suited for underpinning program logics and checking the reachability of an erroneous configuration and/or robustness for finite-state programs with loops (e.g. [Abdulla et al. 2021; Bouajjani et al. 2013; Lahav and Boker 2020]).

Third, the coreRMA authors have failed to *validate* their model against existing implementations in that they could not observe *any* of the weak behaviours allowed by coreRMA on existing implementations. That is, they could not practically justify the weakness of coreRMA.

Fourth and most importantly, as we discuss in detail in §6, coreRMA is not faithful to the RDMA specification [IBTA 2022] and departs from it in three different ways. In particular, coreRMA is neither stronger nor weaker than the specification, meaning that it admits certain behaviours disallowed by the specification, while prohibiting others allowed by the specification.

To close these gaps, we present RDMA$^{\text{TSO}}$, the *first* formal semantics of RDMA programs in the context of the x86 architecture, which implements the TSO model [Sewell et al. 2010]. We describe RDMA$^{\text{TSO}}$ both *operationally* and *declaratively* and prove that the two are *equivalent*. Specifically, we first develop *two operational models* of RDMA$^{\text{TSO}}$: (1) a *concrete* model, reflecting the hardware structure for propagating data across the network; and (2) a *simplified* model, abstracting away the hardware details, resulting in a cleaner model. We prove that our two operational characterisations of RDMA$^{\text{TSO}}$ are equivalent and mechanise our proof in Coq. We then present a declarative model of RDMA$^{\text{TSO}}$ and show that it is equivalent to our simplified (and thus also concrete) operational model.

We have developed RDMA$^{\text{TSO}}$ in close collaboration with engineers at NVIDIA, the largest manufacturer of networking products worldwide (after acquiring Mellanox in 2019). In particular, we have discussed all weak behaviours admitted by RDMA$^{\text{TSO}}$ with the engineers and have reflected the hardware justification for such behaviours in our concrete semantics. To further increase confidence in the fidelity of RDMA$^{\text{TSO}}$ to the specification, we have *empirically validated* it via extensive testing on existing implementations. More specifically, through our empirical validation we have managed to establish that (1) RDMA$^{\text{TSO}}$ is *not too strong*: we did not observe any of the behaviours prohibited by RDMA$^{\text{TSO}}$ on existing implementations; and (2) RDMA$^{\text{TSO}}$ is not *too weak*: we managed to observe *almost all* weak behaviours allowed by RDMA$^{\text{TSO}}$, on existing implementations, and in the few cases where we did not observe a weak behaviour allowed by RDMA$^{\text{TSO}}$, the engineers at NVIDIA

$x = y = 0$	
$x := 1$	$y := 1$
$a := y$	$b := x$

(a) $a = b = 0$ ✓

$x = y = 0$	
$x := 1$	$y := 1$
`mfence`	`mfence`
$a := y$	$b := x$

(b) $a = b = 0$ ✗

$x = y = 0$	
$a := y$	$b := x$
$x := 1$	$y := 1$

(c) $a = b = 1$ ✗

$x = y = 0$	
$x := 1$	$a := y$
$y := 1$	$b := x$

(d) $a = 1, b = 0$ ✗

Fig. 1. TSO litmus tests for CPU concurrency, where locations x, y are accessed by all threads, while locations a, b are accessed by one thread only, and $x = y = 0$ on the first line denotes that x, y initially hold value 0.

confirmed that current implementations explicitly do not utilise the weakness admitted by the RDMA specification (see §5).

Contributions and Outline. In §2 we present an intuitive account of RDMA$^{\text{TSO}}$ through a number of examples. In §3 we present our concrete and simplified operational semantics of RDMA$^{\text{TSO}}$ and show that they are equivalent. In §4 we present our declarative semantics of RDMA$^{\text{TSO}}$ and show that it is equivalent to our simplified operational semantics. In §5 we describe how we empirically validated RDMA$^{\text{TSO}}$ through extensive litmus testing. We discuss related and future work in §6.

Additional Material. The full proofs of all stated theorems are given in the accompanying appendix and Coq development [Ambal et al. 2024]. We provide the executable RDMA code (in machine-readable format) and detailed instructions for replicating our experiments and analysing our litmus tests [Ambal et al. 2024].

2 OVERVIEW

We present an account of the formal RDMA semantics, RDMA$^{\text{TSO}}$, through a number of examples. We model concurrent programs running over a network of machines, and hereafter refer to each machine on the network as a *node*. In our setting, the semantics of a concurrent program and thus its possible weak behaviours are determined by two factors: (1) the *origin* of the threads, i.e. whether all threads originate from (are forked by) the same node and thus the concurrency is *intra-node* (within one node), or they originate from several nodes and thus concurrency is *inter-node* (across two or mode nodes); and (2) the *memory* targeted by the threads, i.e. whether each thread accesses its own local memory (on the same node), that of other nodes, or a combination thereof.

Litmus Test Outcome Notation. In the remainder of this article, as well as in the technical appendix, we present small representative examples (known as litmus tests in the literature) to illustrate whether an outcome is allowed by a given model (e.g. in Fig. 1, Fig. 2 and Fig. 3), and annotate a given outcomes with: (1) ✓, to denote that the outcome is *allowed* by the model *and observed* in practice (in our empirical validation); (2) ✓*, to denote that the outcome is *allowed* by the model *and not observable* in practice; (3) ✗, to denote that the outcome is *disallowed* by the model *and not observed* in practice. See §5 or §A for more details.

CPU Concurrency and TSO. Existing work on RDMA semantics [Dan et al. 2016] assumes that CPU concurrency on each node is governed by the strong and unrealistic *sequential consistency* (SC) model [Lamport 1979]. We relax this assumption here and instead model each node as an x86 machine, subject to the TSO memory model [Sewell et al. 2010] introduced by the SPARC architecture [SPARC 1992]. Under TSO, a later read (in program order) can be reordered before an earlier write on a different location. This is illustrated in the *store buffering* example of Fig. 1a, where x and y denote locations accessed by both threads, while a and b denote locations accessed by

$x=0$	$z=0$
$x := 1$ $z^2 := x$	

(a) $z=0$ ✗ $z=1$ ✓

$x=0$	$z=0$
$z^2 := x$ $x := 1$	

(b) $z=0$ ✓ $z=1$ ✓

$x=0$	$z=0$
$z^2 := x$ `poll(2)` $x := 1$	

(c) $z=0$ ✓ $z = 1$ ✗

$x=0$	$z=0$
$z^2 := x$ $z^2 := x$ `poll(2)` $x := 1$	

(d) $z=0$ ✓ $z=1$ ✓

$x=0$	$z=1$	$y=2$
$x := z^2$ $x := y^3$		

(e) $x=1$ ✓ $x=2$ ✓

$x=0$	$z=1$	$y=2$
$y^3 := x$ $x := z^2$		

(f) $y=0$ ✓ $y=1$ ✓

$x=1$	$y=z=0$
$z^2 := x$ $x := y^2$	

(g) $z=0$ ✗ $z=1$ ✓

$x = 1$	$y=z=0$
$x := y^2$ $z^2 := x$	

(h) $z=0$ ✓ $z=1$ ✓

$x = 1$	$y=z=0$
$x := y^2$ `rfence (2)` $z^2 := x$	

(i) $z=0$ ✓ $z=1$ ✗

Fig. 2. Sequential RDMA litmus tests (excerpt), where each column (separated by ||) denotes a distinct node, the statement on the top line of each column denotes the initial values of locations, and the statements in the caption express whether each outcome is allowed by RDMA$^{\text{TSO}}$ and observed in practice (✓), or disallowed by RDMA$^{\text{TSO}}$ and not observed in practice (✗); i.e. we have empirically validated all outcomes shown.

one thread only.[1] Specifically, the reads $a := y$ and $b := x$ can respectively be reordered before the writes $x := 1$ and $y := 1$, allowing them to read the initial value 0 from y and x, yielding $a = b = 0$ at the end of execution. Note that this weak behaviour is not allowed under the stronger SC model as SC admits no instruction reordering.

To prevent such write-read reordering, one can use an `mfence` as in Fig. 1b: `mfence` instructions cannot be reordered in either direction, and thus the weak behaviour shown is no longer possible. Indeed, other than write-read reordering, TSO admits no other reorderings and thus other weak behaviours, e.g. *load buffering* in Fig. 1c and *message passing* in Fig. 1d are prohibited under TSO.

Remote Direct Memory Access (RDMA). RDMA allows one to build a network of communicating nodes (machines), where each node can *directly* access remote memory (of other nodes) through its network interface card (NIC). RDMA networks are programmed via operations that read from and write to remote memory, as well as various synchronisation operations. As such, programming RDMA networks is conceptually similar to shared memory systems such as TSO.

To distinguish remote (RDMA) operations from CPU ones, we refer to RDMA reads and writes as *get* and *put* operations, respectively. Moreover, to distinguish local and remote memory locations, we write x^n for a memory location on a remote node n, and write x for a memory location on the current local node. A put operation is of the form $x^n := y$, and consists of reading from a local memory location y (referred to as a 'NIC local read') and writing to a remote memory location x on node n (a 'NIC remote write' or simply a 'remote write'). Similarly, a get operation is of the form $x := y^n$, and consists of reading from a remote memory location y on node n (a 'NIC remote read' or a 'remote read') and writing to a local memory location x (a 'NIC local write'). We write $\overline{n}$ to identify a node other than n. When a thread on local node n issues a remote operation to be executed on remote node $\overline{n}$, we denote this by stating that the operation is *by n towards* $\overline{n}$.

Sequential (Single-Threaded) RDMA$^{\text{TSO}}$ Behaviours. When a thread issues a put or get operation, it is handled by the NIC subsystem (and its associated queue pairs and buffers as shown

[1] In our general model, all memory locations are shared and thus can be accessed by all threads both locally (on the same node) and remotely. However, for better readability, we follow the convention of naming locations accessed by multiple threads (locally or remotely) as x, y, z and w, while naming locations accessed by a single local thread as a, b, c and d.

in Fig. 4), in contrast to the CPU operations which are handled by the processor subsystem (and its associated store buffers). As such, the interaction between CPU and remote operations lead to further behaviours even within a *sequential* (single-threaded) program. We demonstrate this in the examples of Fig. 2, where each column represents a distinct node, numbered from 1 onwards. For instance, the example in Fig. 2a comprises a single thread on node 1 (the left-most column) that writes to the local location x ($x := 1$) and puts x towards the remote location z on node 2 ($z^2 := x$).

Intuitively, when a thread t on n issues remote operations towards node $\overline{n}$, one can view these remote operations as if being executed by a thread running *in parallel* to t. As such, when a remote operation *follows* a CPU one, the order of the two operations is preserved since the parallel thread is spawned only after the CPU operation is executed. This is illustrated in Fig. 2a: as $z^2 := x$ follows $x := 1$, it observes the $x := 1$ write and thus puts value 1 to z; i.e. outcome $z = 0$ is not permitted.

By contrast, when a remote operation *precedes* a CPU one, the remote operation is performed by a 'separate thread' run in parallel to the later CPU operation in the main thread, and thus may execute before or after the CPU operation, meaning that in the latter case the execution order is not preserved. This is illustrated in Fig. 2b, where the earlier $z^2 := x$ may execute (be reordered) after the later $x := 1$, and thus both $z = 0$ and $z = 1$ outcomes are possible.

Therefore, before using the result of a get or reusing the memory location of a put, it may be desirable to avoid such reordering and to wait for the remote operation to complete. This can be done through a CPU *poll* operation, $\texttt{poll}(n)$, that blocks until the *earliest* (in program order) remote operation towards node n has completed.[2] This is shown in Fig. 2c, obtained from Fig. 2b by inserting a poll after the remote operation: $\texttt{poll}(2)$ waits for $z^2 := x$ to complete before proceeding with $x := 1$, and thus $z^2 := x$ can no longer be reordered after $x := 1$, prohibiting the $z = 1$ outcome.

Note that each $\texttt{poll}(n)$ waits for *only one* (the earliest in program order) and *not all* pending remote operation towards n to complete. This allows for more efficient and fine-grained control over remote operations, but requires some care. For instance, consider the example in Fig. 2d, where $\texttt{poll}(2)$ blocks until the *first* $z^2 := x$ is complete, and thus the second $z^2 := x$ operation can be reordered after the later $x := 1$, once again allowing for the $z = 1$ outcome.

Two remote operations towards *different* nodes are fully independent and can execute in either order (as if within two separate threads). For instance, the two get operations in Fig. 2e can execute in either order, and thus the final value of x may be either that of z ($x = 1$ outcome) or that of y (the $x = 2$ outcome). Similarly, Fig. 2f has two possible outcomes. The only way to enforce the execution order is polling the first remote operation before running the second.

The ordering guarantees on remote operations towards the *same* node are stronger and only certain reorderings are allowed. Recall that a remote (NIC) put operation $x^n := y$ comprises two steps: a NIC local read (obtaining the value of y) and a NIC remote write (writing the value of y to x^n). Similarly, a remote get operation $x := y^n$ comprises two steps: a NIC remote read (obtaining the value of y^n) and a NIC local write (writing the value of y^n to x). Intuitively, for remote operations on a given location x, these four steps mandate a *precedence* which in turn determines whether they can be reordered. Specifically, the four steps described above give way to the following precedence order: i) NIC local read; ii) NIC remote write; iii) NIC remote read; iv) NIC local write.

If a step with a higher precedence (e.g. a NIC local read) is in program order before one with a lower precedence (e.g. a NIC local write), then their order is preserved and they cannot be reordered; otherwise the order is not necessarily preserved and these steps can be reordered. For instance, in the Fig. 2g example, the earlier NIC local read on x (in $z^2 := x$) has a higher precedence than the

[2] In the Verbs API, $\texttt{poll}(n)$ returns a value denoting whether the operation being polled has completed and does not block the execution. However, it is common practice for calls to poll to be placed in a spin loop that returns only when the operation is completed, hence blocking the execution. Here, we model this common pattern with stronger behaviours.

$y=0$	$x=0$
$x^2 := 1$	$y^1 := 1$
$a := y$	$b := x$

(a) $a=b=0$ ✓

$x=0$	$y=0$
$a := y^2$	$b := x^1$
$x := 1$	$y := 1$

(b) $a=b=1$ ✓

$x=0$	$y=0$
$a := y^2$	$b := x^1$
`poll(2)`	`poll(1)`
$x := 1$	$y := 1$

(c) $a=b=1$ ✗

$y=0$	$x=0$
$x^2 := 1$	$y^1 := 1$
`poll(2)`	`poll(1)`
$a := y$	$b := x$

(d) $a=b=0$ ✓*

$y=w=0$	$x=z=0$
$x^2 := 1$	$y^1 := 1$
$c := z^2$	$d := w^1$
`poll(2)`	`poll(1)`
`poll(2)`	`poll(1)`
$a := y$	$b := x$

(e) $a=b=0$ ✗

Fig. 3. Concurrent RDMA litmus tests (excerpt). The annotations in the captions indicate the given outcome is allowed by RDMA$^{\text{TSO}}$ *and observed* in our empirical validation (✓), allowed by RDMA$^{\text{TSO}}$ *and not observable* in practice (✓*), or disallowed by RDMA$^{\text{TSO}}$ *and not observed* in our validation (✗). That is, we have empirically validated all outcomes shown, except that in (d), which is due to the underlying implementation explicitly not utilising the weakness admitted by the specification – see §5 or §A for more details.

later NIC local write on x (in $x := y^2$), and thus the order of these steps on x is preserved; i.e. the old value (1) of x is written to z^2 leading to the $z = 1$ outcome, and the $z = 0$ outcome is disallowed.

By contrast, in the Fig. 2h example, the earlier NIC local write on x (in $x := y^2$) has a lower precedence than the later NIC local read on x (in $z^2 := x$) and thus the two steps can be reordered. Besides the SC outcome ($z = 0$), the program might execute the NIC local read on x before the NIC local write on x, thereby reading the initial value 1 from x and writing it to z (outcome $z = 1$).

As before, the reordering of the two remote operations in Fig. 2h can be prevented by polling the first operation before the second operation. However, polling is costly as it leads to global blocking: it blocks the current thread *both* on the CPU and the NIC towards *all* nodes (i.e. the current thread cannot execute any remote operations on *any* node). Alternatively, we can use a *remote fence*[3]operation, `rfence` (n), that blocks the current thread *only* on the NIC and *only* towards node n. (i.e. the thread cannot execute any remote operations towards n, but can execute both on the CPU as well as on the NIC towards nodes other than n). This in turn ensures that earlier (in program order) remote operations by the thread towards n (those before the fence) are executed before later remote operations towards n (those after the fence). This is illustrated in Fig. 2i, obtained from Fig. 2h by inserting `rfence` (2) between the two remote operations towards node 2, thereby ensuring that they are executed in order, and thus $z = 1$ is no longer possible.

Concurrent (Multi-threaded) RDMA$^{\text{TSO}}$ Behaviours. The real power of RDMA comes from multiple programs running on different nodes. This introduces a wide range of weak behaviours, as we describe below. A network can comprise several nodes, each running several concurrent threads. Here, we focus on a few examples each with two nodes, with each node comprising a single thread.

The store buffering behaviour due to the TSO model discussed in Fig. 1a is also possible in the RDMA setting, i.e. when locations x and y are on two different nodes, as shown in Fig. 3a. Moreover, further weak behaviours not possible under TSO, e.g. load buffering in Fig. 1c, are permitted in the RDMA setting, as shown in Fig. 3b. As before, this can be intuitively justified by conceptually viewing the remote operations in each thread as being executed by a separately spawned thread.

Most weak behaviours such as load buffering in Fig. 3b can be prevented by polling the remote operations as needed, as shown in Fig. 3c. Specifically, the poll operations in Fig. 3c await the completion of the preceding get operations, and thus the earlier get operations cannot be reordered after the later writes, thus prohibiting the weak behaviour $a=b=1$. However, a notable exception to this is the store buffering weak behaviour which cannot be prevented even when polling the

[3] For InfiniBand Verbs, this remote fence is not an independent operation but a flag that can be set on the later operation.

remote operations, as shown in Fig. 3d. This is because the specification of polling offers different guarantees for get and put operations. Specifically, polling a get operation $a := x^2$ offers a strong guarantee and behaves intuitively: polling ensures that the get operation is complete (i.e. the value of x^2 is fetched from node 2), and the executing thread performs the associated NIC local write on a before marking the operation as complete and proceeding with the remainder of the execution. By contrast, polling a put operation $x^2 := 1$ offers a weaker guarantee: when sending value 1 towards node 2 to be put in x^2, the remote NIC merely *acknowledges* having received the data (value 1), but this data may still reside in a *buffer* (i.e. the PCIe fabric) of the remote node, *pending* to be written to the remote memory. Polling a put operation then simply awaits the acknowledgement of the data receipt and not its full completion (the data being written to memory). As such, it is possible to poll a put operation successfully before the associated remote write has fully completed. In the case of store buffering in Fig. 3d, it is possible for both poll operations to complete before the values of x and y are updated (to 1) in memory, and thus for the later reads to read their old values (0). However, note that existing implementations on current hardware (including ones against which we validated our model,) do not utilise this flexibility admitted by the specification. As such, the weak behaviour in Fig. 3d is not observable in practice, and thus we could not observe them in our validation effort (as indicated by ✓*) – see §5. Nevertheless, since our aim is to capture the *specification* and not the *implementation*, we have modelled RDMA$^{\text{TSO}}$ to allow this weak behaviour.

The behaviours discussed thus far all hold of the general RDMA specification [IBTA 2022] as well as the PCIe standard [PCI-SIG 2022]. However, in certain cases the PCIe standard offers stronger guarantees than those delineated by the RDMA specification. In particular, PCIe does not allow a read to fetch a pending value that has not yet been committed to memory. As such, a NIC remote read *flushes* (commits) all pending NIC remote writes to memory, while a NIC local read flushes all pending NIC local writes to memory. Interestingly, we can use this guarantee to prevent weak behaviours such as store buffering (which, in theory, cannot be prevented even via polling). Specifically, recall that polling a put only ensures that the data transmitted has reached the remote node and may not have yet been committed to its memory. However, by polling a later get (towards the same remote node) we can ensure the previous NIC remote writes (including that of the recently polled put) have been committed to the remote memory. An example of this is shown in Fig. 3e, obtained from Fig. 3d by adding additional gets ($c := z^2$ and $d := w^1$) and subsequently polling them. Reading from z^2 and w^2 in effect flushes the pending NIC remote writes on both nodes, ensuring that the effects of the earlier puts ($x^2 := 1$ and $y^1 := 1$) are committed to memory, which in turn ensures that the later $a := y$ and $b := x$ reads observe the updates on y and x (value 1), thus prohibiting $a = b = 0$.

As mentioned above, this guarantee is PCIe-specific, and not mentioned by the RDMA standard. However, as PCIe is the *de facto* standard for RDMA programming, and since all widely available RDMA hardware is PCIe-compatible, here we opt to model this guarantee, resulting in a stronger model. Nevertheless, in §3 and §4 we also describe how we can weaken our models by removing this guarantee, both in our operational and declarative semantics.

Our aim here was to provide an overview of the weak RDMA behaviours both in the sequential and concurrent settings. We refer the reader to §A for further examples of weak behaviours.

3 RDMA$^{\text{TSO}}$ CONCRETE AND SIMPLIFIED OPERATIONAL SEMANTICS

We begin with several preliminary concepts. We present an informal account of our concrete semantics (§3.1), our formal concrete semantics (§3.2) and our equivalent simplified semantics (§3.3).

Nodes and Threads. We consider a system with N nodes (machines) and M threads in total across all machines. Let $\text{Node} = \{1, \ldots, N\}$ be the set of *node identifiers*, and $\text{Tid} = \{1, \ldots, M\}$ be the set of *thread identifiers*. We use n and t to range over Node and Tid, respectively. Given a node n, we write $\overline{n}$ to range over $\mathsf{Node} \setminus \{n\}$. Each thread $t \in \mathsf{Tid}$ is associated with a node, written $n(t)$. Note that multiple threads may run on the same node.

Memory Locations. Each node n has a set of *locations*, Loc_n, accessible by all nodes. We define $\mathsf{Loc} \triangleq \biguplus_n \mathsf{Loc}_n$ and $\mathsf{Loc}_{\overline{n}} \triangleq \mathsf{Loc} \setminus \mathsf{Loc}_n$. We use x^n, y^n, z^n and $x^{\overline{n}}, y^{\overline{n}}, z^{\overline{n}}$ to range over Loc_n and $\mathsf{Loc}_{\overline{n}}$, respectively. When the choice of n is clear, we write x for x^n; similarly for $\overline{n}$. For clarity, we use distinct location names across nodes, and thus write $n(x)$ for the unique $n \in \mathsf{Node}$ where $x \in \mathsf{Loc}_n$.

Values and Expressions. We assume a set of values, Val, with $\mathbb{N} \subseteq \mathsf{Val}$, and use v to range over Val. We assume a language of expressions over Val and Loc, and elide its exact syntax and semantics (as these are standard). We denote by Exp the set of all expressions and use e to range over Exp. We write $[[e]]$ for the semantic evaluation of a *closed* expression e (i.e. one without any locations), $\mathsf{elocs}(e)$ for the locations in an expression e, and $e[v/x]$ for the expression obtained from e after substituting all occurrences of x by v. We use e^n to range over expressions with $\mathsf{elocs}(e^n) \subseteq \mathsf{Loc}_n$.

Sequential Commands and Programs. *Sequential* programs running on node n are described by the C^n grammar below and include primitive commands (c^n), sequential composition ($\mathsf{C}^n_1; \mathsf{C}^n_2$), non-deterministic choice ($\mathsf{C}^n_1 + \mathsf{C}^n_2$, executing either C^n_1 or C^n_2) and non-deterministic loops (C^{n*}, executing C^n for a finite, possibly zero, number of iterations). A (concurrent) *program*, P, is a map from thread identifiers to commands, associating each thread $t \in \mathsf{Tid}$ with a command on node $n(t)$.

$$\mathsf{Comm} \ni \mathsf{C}^n ::= \texttt{skip} \mid \mathsf{c}^n \mid \mathsf{C}^n_1; \mathsf{C}^n_2 \mid \mathsf{C}^n_1 + \mathsf{C}^n_2 \mid \mathsf{C}^{n*} \qquad \mathsf{PComm} \ni \mathsf{c}^n ::= \mathsf{cc}^n \mid \mathsf{rc}^n$$

$$\mathsf{CComm} \ni \mathsf{cc}^n ::= x := e^n \mid \texttt{assume}(x = v) \mid \texttt{assume}(x \neq v) \mid \texttt{mfence} \mid x := \texttt{CAS}(y, e_1, e_2) \mid \texttt{poll}\ (\overline{n})$$

$$\mathsf{RComm} \ni \mathsf{rc}^n ::= x := \overline{y} \mid \overline{y} := x \mid \texttt{rfence}\ (\overline{n})$$

Primitive commands include *CPU* (cc^n) and *RDMA* (rc^n) operations. A CPU operation on n may be a no-op (`skip`), an assignment to a local location ($x := e$), an assumption on the value of a local location ($\texttt{assume}(x = v)$ and $\texttt{assume}(x \neq v)$), a memory fence (`mfence`), an atomic CAS ('compare-and-set') operation ($x := \texttt{CAS}(y, e_1, e_2)$), or a 'poll', $\texttt{poll}(\overline{n})$, that awaits the completion notification of the earliest put/get that is pending (not yet acknowledged). An RDMA operation may be (i) a 'get', $x := \overline{y}$, reading from remote location $\overline{y}$ and writing the result to local location x; (ii) a 'put', $\overline{y} := x$, reading from local location x and writing the result to remote location $\overline{y}$; or (iii) an 'RDMA fence', $\texttt{rfence}\ (\overline{n})$, which ensures that all later (in program order) RDMA operations towards $\overline{n}$ will await the completion of all earlier RDMA operations towards $\overline{n}$. Note that $\texttt{poll}\ (\overline{n})$ is executed by the CPU and blocks its thread (and prevents the requests of later remote operations), while $\texttt{rfence}\ (\overline{n})$ blocks the NIC for the execution of remote operations towards node $\overline{n}$. In what follows, we write $\overline{x} := 1$ as a shorthand for $\overline{x} := a$ for some local location a containing value 1.

3.1 RDMA$^{\text{TSO}}$ Concrete Operational Semantics at a Glance

RDMA$^{\text{TSO}}$ Architecture and CPU Operations. Conceptually, the RDMA$^{\text{TSO}}$ architecture is as shown at the top-right of Fig. 4, where for brevity we depict two nodes, each comprising (1) m threads and their associated FIFO (first-in-first-out) *store buffers*; (2) a network interface card (NIC) and its PCIe root complex; and (3) the memory. Store buffers are used to model write-read reordering on TSO, which account for the weak behaviour in Fig. 1a. Specifically, executing a write on TSO comprises two steps: (i) when a thread issues a write, the write is only recorded in its store buffer; (ii) writes in the buffer are debuffered (in FIFO order) and propagated to the local memory at a later point. When a thread issues a read from a location x, it first consults its own store buffer. If it contains a write for x, the thread reads the value of the latest such write; otherwise, the thread reads the value of x from its local memory. In other words, one can model the reordering of a write w

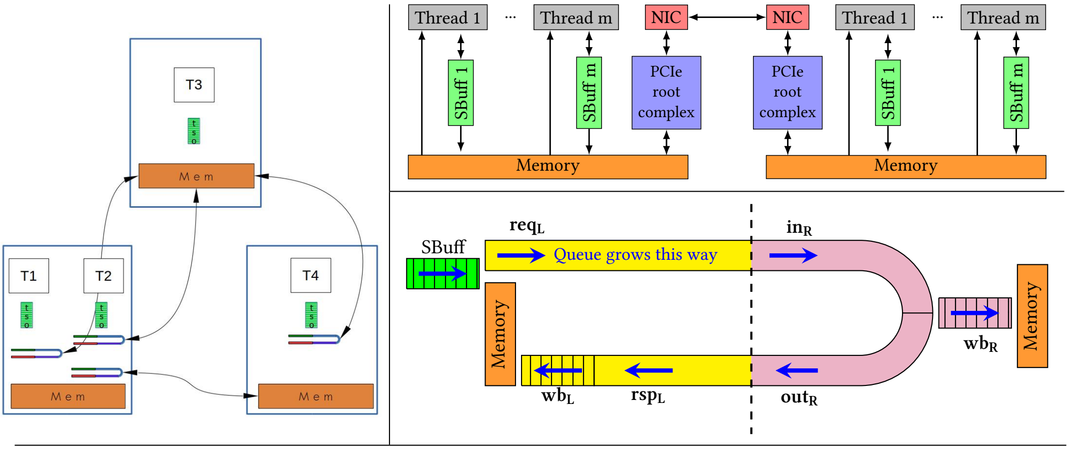

Fig. 4. RDMA$^{\text{TSO}}$ architecture overview (top, right); a possible RDMA network configuration with three nodes and four thread (left); the queue-pair structure (below, right)

	$y^{\overline{n}} := x^n$	$y^{\overline{n}} := v$	ack_p	$x^n := y^{\overline{n}}$	$x^n := v$	cn	rfence $\overline{n}$
b ∈ SBuff	✓			✓	✓		✓
$\mathbf{req_L}$	✓			✓			✓
$\mathbf{in_R}$		✓		✓			
$\mathbf{wb_R}$		✓					
$\mathbf{out_R}$			✓	✓	✓		
$\mathbf{rsp_L}$			✓		✓		
$\mathbf{wb_L}$					✓	✓	

Fig. 5. The types of entries in the store buffers (b) of a thread on n and its six buffers on the $\overline{n}$-queue pair

after a later read r by *delaying* the debuffering of w until after r has executed. Moreover, executing an `mfence` or a CAS debuffers all its delayed writes in the store buffer and propagates them to memory (in FIFO order), thus preventing write-read reordering. That is, the only *CPU* operations that a store buffer contains are delayed writes. We describe the execution of polls shortly.

Remote Operations and Queue Pairs. To model the communication between the nodes, each thread t has a distinct *queue pair* for each remote node whose memory t accesses. For instance, the network configuration in the left of Fig. 4 comprises three nodes with four threads and with the queue pairs depicted as 'horse-shoes', where e.g. thread T2 in the bottom left node accesses the memories of the other two nodes, and it does so through two distinct queue pairs shown underneath T2. We refer to the queue pair of a thread towards node $\overline{n}$ as its $\overline{n}$-queue pair.

The details of a queue-pair structure is shown at the bottom-right of Fig. 4, where each queue pair of thread t is connected to the store buffer of t. As shown, an $\overline{n}$-queue pair comprises six FIFO buffers—$\mathbf{req_L}$, $\mathbf{in_R}$, $\mathbf{wb_R}$, $\mathbf{out_R}$, $\mathbf{rsp_L}$, and $\mathbf{wb_L}$—that we describe below. A queue pair contains *pending* operations if any of its $\mathbf{in_R}$, $\mathbf{out_R}$ or $\mathbf{rsp_L}$ components is non-empty. The types of the entries in each of the six buffers of a queue pair is also summarised in Fig. 5.

Recall that a remote operation rc^n may either be a get, put or remote fence instruction. When t executes a remote operation rc^n towards node $\overline{n}$, it adds rc^n to its store buffer. That is, as well as CPU (delayed) writes, a store buffer may contain remote (get, put and fence) operations. Once rc^n is debuffered from the store buffer, it is forwarded to its $\overline{n}$-queue pair, where it travels through the queue-pair pipeline and is processed differently depending on the type of rc^n, as we describe below.

Request Buffers ($\mathbf{req}_\mathrm{L}$). The $\mathbf{req}_\mathrm{L}$ is the entry point of the queue pair, containing remote requests that are to be forwarded. It comprises a sequence of remote get, put and fence operations to be handled by the local NIC. When an entry rc^n reaches the head of $\mathbf{req}_\mathrm{L}$, it is processed as follows. (1) If rc^n is a get, then it is simply forwarded to the remote inbox $\mathbf{in}_\mathrm{R}$. (2) If rc^n is a put $x^{\overline{n}} := y$, then the value v of y is looked up in the local memory and $x^{\overline{n}} := v$ is forwarded to $\mathbf{in}_\mathrm{R}$; i.e. y in $x^{\overline{n}} := y$ is replaced with its current in-memory value. (3) If rc^n is a remote fence, then the execution on the queue pair is blocked until it has no pending operations, i.e. $\mathbf{in}_\mathrm{R}$, $\mathbf{out}_\mathrm{R}$ and $\mathbf{rsp}_\mathrm{L}$ are all empty.

Inbox Buffers ($\mathbf{in}_\mathrm{R}$). The $\mathbf{in}_\mathrm{R}$ contains requests forwarded by $\mathbf{req}_\mathrm{L}$ that are to be processed by $\overline{n}$, i.e. each $\mathbf{req}_\mathrm{L}$ entry is either a get or a put with a value (of the form $y^{\overline{n}} := v$). Processing a put $y^{\overline{n}} := v$ (once at $\mathbf{in}_\mathrm{R}$ head) forwards it to the remote write-back buffer $\mathbf{wb}_\mathrm{R}$ and also sends a *put acknowledgement*, ack_p, to $\mathbf{out}_\mathrm{R}$. Processing a get $x := y^{\overline{n}}$ (once at $\mathbf{in}_\mathrm{R}$ head) does not immediately fetch the y value from the $\overline{n}$ memory; rather, it forwards it to $\mathbf{out}_\mathrm{R}$ to be fulfilled later in $\mathbf{out}_\mathrm{R}$.

Outbox Buffers ($\mathbf{out}_\mathrm{R}$). The $\mathbf{out}_\mathrm{R}$ buffer contains requests processed or forwarded by $\mathbf{in}_\mathrm{R}$, and thus each entry in $\mathbf{out}_\mathrm{R}$ is either a put acknowledgement (ack_p), a get operation $x := y^{\overline{n}}$ *yet to be fulfilled*, or a *fulfilled* get $x := v$. An acknowledgement or a fulfilled get is processed when it reaches the head of $\mathbf{out}_\mathrm{R}$, whereupon it is simply forwarded to $\mathbf{rsp}_\mathrm{L}$. By contrast, a yet-to-be-fulfilled get $x := y^{\overline{n}}$ may be fulfilled at *any time* (before reaching the $\mathbf{out}_\mathrm{R}$ head), where the value v of y is fetched from the memory of $\overline{n}$, and $x := y^{\overline{n}}$ is transformed to the fulfilled get $x := v$ and left in $\mathbf{out}_\mathrm{R}$. This fulfilled get is later processed once it reaches the head of $\mathbf{out}_\mathrm{R}$, as described above.

Remote Write-Back Buffers ($\mathbf{wb}_\mathrm{R}$). The $\mathbf{wb}_\mathrm{R}$ buffer contains requests forwarded by $\mathbf{in}_\mathrm{R}$, and thus each entry in $\mathbf{wb}_\mathrm{R}$ is a put operation with a value (of the form $y^{\overline{n}} := v$). Processing $y^{\overline{n}} := v$ (at the head of $\mathbf{wb}_\mathrm{R}$) simply removes it from $\mathbf{wb}_\mathrm{R}$ and writes v to y in the memory of $\overline{n}$.

Response Buffers ($\mathbf{rsp}_\mathrm{L}$). The $\mathbf{rsp}_\mathrm{L}$ buffer contains processed requests forwarded by $\mathbf{out}_\mathrm{R}$, and thus each entry in $\mathbf{rsp}_\mathrm{L}$ is either an ack_p, acknowledging a processed put, or a fulfilled get. Processing each $\mathbf{rsp}_\mathrm{L}$ entry (once at the head of $\mathbf{rsp}_\mathrm{L}$) generates a *completion notification*, cn, which is used to serve poll requests, as we describe shortly. Specifically, processing ack_p simply removes it from $\mathbf{rsp}_\mathrm{L}$ and forwards a completion notification to $\mathbf{wb}_\mathrm{L}$. Analogously, processing a fulfilled get $x := v$ simply forwards $x := v$ together with a completion notification to $\mathbf{wb}_\mathrm{L}$.

Local Write-Back Buffers ($\mathbf{wb}_\mathrm{L}$). The $\mathbf{wb}_\mathrm{L}$ buffer contains processed requests forwarded by $\mathbf{rsp}_\mathrm{L}$; i.e. each entry in $\mathbf{wb}_\mathrm{L}$ is either a completion notification (cn, associated with processed gets and puts), or a fulfilled get (of the form $x := v$). Processing a fulfilled get $x := v$ simply removes it and writes v to x in the local memory. A completion notification is left in the $\mathbf{wb}_\mathrm{L}$ and is only removed when the associated get/put operation is polled, as we describe below.[4]

Poll Operations. Once a get enters $\mathbf{req}_\mathrm{L}$, it progresses through the pipeline as follows. (G_1) it is forwarded to $\mathbf{in}_\mathrm{R}$; (G_2) it is forwarded to $\mathbf{out}_\mathrm{R}$ without being fulfilled; (G_3) it is fulfilled at some point while in $\mathbf{out}_\mathrm{R}$; (G_4) it is forwarded to $\mathbf{rsp}_\mathrm{L}$; (G_5) it is forwarded to $\mathbf{wb}_\mathrm{L}$ together with a completion notification (cn); and (G_6) it is removed from $\mathbf{wb}_\mathrm{L}$ and its effect is written to memory.

Similarly, once a put operation $x^{\overline{n}} := y$ enters the queue-pair pipeline it proceeds as follows. (P_1) it is simplified to $x^{\overline{n}} := v$ (where v is the value of y in the local memory) and forwarded to $\mathbf{in}_\mathrm{R}$; (P_2) it is forwarded to $\mathbf{wb}_\mathrm{R}$ and simultaneously an acknowledgement ack_p is forwarded to $\mathbf{out}_\mathrm{R}$; (P_3) its $\mathbf{wb}_\mathrm{R}$ entry is eventually removed and applied to the remote memory; its associated ack_p in $\mathbf{out}_\mathrm{R}$ (P_4) is forwarded to $\mathbf{rsp}_\mathrm{L}$; and (P_5) later forwarded to $\mathbf{wb}_\mathrm{L}$ as cn.

[4] In some implementations, write-back buffers ($\mathbf{wb}_\mathrm{L}$ and $\mathbf{wb}_\mathrm{R}$) of different queue pairs may physically use the same hardware buffers in the PCIe fabric. This does not introduce any additional weak behaviours.

Program transitions: $\mathsf{Prog} \xrightarrow{\mathsf{Tid}:\mathsf{Lab} \uplus \{\varepsilon\}} \mathsf{Prog}$ **Command transitions:** $\mathsf{Comm} \xrightarrow{\mathsf{Lab} \uplus \{\varepsilon\}} \mathsf{Comm}$

$$\mathsf{Lab} \triangleq \bigcup_{n} \mathsf{Lab}_n \qquad l \in \mathsf{Lab}_n \triangleq \left\{ \begin{matrix} \mathtt{lW}(x^n, v), \mathtt{lR}(x^n, v), \mathtt{CASS}(x^n, v_1, v_2), \mathtt{CASF}(x^n, v), \\ \mathtt{F}, \mathtt{P}(\overline{n}), \mathtt{Get}(x^n, y^{\overline{n}}), \mathtt{Put}(y^{\overline{n}}, x^n), \mathtt{rF}(\overline{n}) \end{matrix} \;\middle|\; \begin{matrix} x, y \in \mathsf{Loc}, \\ v, v_1, v_2 \in \mathsf{Val} \end{matrix} \right\}$$

$$\frac{\mathsf{C}_1 \xrightarrow{l} \mathsf{C}_1'}{\mathsf{C}_1;\mathsf{C}_2 \xrightarrow{l} \mathsf{C}_1';\mathsf{C}_2} \qquad \frac{}{\mathtt{skip};\mathsf{C} \xrightarrow{\varepsilon} \mathsf{C}} \qquad \frac{i \in \{1,2\}}{\mathsf{C}_1 + \mathsf{C}_2 \xrightarrow{\varepsilon} \mathsf{C}_i^n} \qquad \frac{}{\mathsf{C}^* \xrightarrow{\varepsilon} \mathtt{skip}} \qquad \frac{}{\mathsf{C}^* \xrightarrow{\varepsilon} \mathsf{C};\mathsf{C}^*}$$

$$\frac{\mathsf{C} \rightsquigarrow \mathsf{C}'}{\mathsf{C} \xrightarrow{\varepsilon} \mathsf{C}'} \qquad \frac{\mathsf{elocs}(e) = \emptyset}{x := e \xrightarrow{\mathtt{lW}(x, [\![e]\!])} \mathtt{skip}} \qquad \frac{\mathsf{elocs}(e_{\mathrm{old}}) = \mathsf{elocs}(e_{\mathrm{new}}) = \emptyset \quad v \neq [\![e_{\mathrm{old}}]\!]}{z := \mathtt{CAS}(x, e_{\mathrm{old}}, e_{\mathrm{new}}) \xrightarrow{\mathtt{CASF}(x, v)} z := v}$$

$$\frac{\mathsf{elocs}(e_{\mathrm{old}}) = \mathsf{elocs}(e_{\mathrm{new}}) = \emptyset}{z := \mathtt{CAS}(x, e_{\mathrm{old}}, e_{\mathrm{new}}) \xrightarrow{\mathtt{CASS}(x, [\![e_{\mathrm{old}}]\!], [\![e_{\mathrm{new}}]\!])} z := [\![e_{\mathrm{old}}]\!]} \qquad \frac{}{\mathtt{mfence} \xrightarrow{\mathtt{F}} \mathtt{skip}} \qquad \frac{}{x := \overline{y} \xrightarrow{\mathtt{Get}(x, \overline{y})} \mathtt{skip}}$$

$$\frac{}{\overline{y} := x \xrightarrow{\mathtt{Put}(\overline{y}, x)} \mathtt{skip}} \qquad \frac{}{\mathtt{rfence}\ \overline{n} \xrightarrow{\mathtt{rF}(\overline{n})} \mathtt{skip}} \qquad \frac{}{\mathtt{poll}(\overline{n}) \xrightarrow{\mathtt{P}(\overline{n})} \mathtt{skip}}$$

$$\frac{}{\mathtt{assume}(x = v) \xrightarrow{\mathtt{lR}(x, v)} \mathtt{skip}} \qquad \frac{v \neq v'}{\mathtt{assume}(x \neq v') \xrightarrow{\mathtt{lR}(x, v)} \mathtt{skip}} \qquad \frac{\mathsf{P}(t) \xrightarrow{l} \mathsf{C}}{\mathsf{P} \xrightarrow{t:l} \mathsf{P}[t \mapsto \mathsf{C}]}$$

$$\begin{array}{ll} x := e \rightsquigarrow \mathtt{assume}(y = v); x := e[v/y] & \text{for } y \in \mathsf{elocs}(e), v \in \mathsf{Val} \\ z := \mathtt{CAS}(x, e_{\mathrm{old}}, e_{\mathrm{new}}) \rightsquigarrow \mathtt{assume}(y = v); z := \mathtt{CAS}(x, e_{\mathrm{old}}[v/y], e_{\mathrm{new}}) & \text{for } y \in \mathsf{elocs}(e_{\mathrm{old}}), v \in \mathsf{Val} \\ z := \mathtt{CAS}(x, e_{\mathrm{old}}, e_{\mathrm{new}}) \rightsquigarrow \mathtt{assume}(y = v); z := \mathtt{CAS}(x, e_{\mathrm{old}}, e_{\mathrm{new}}[v/y]) & \text{for } y \in \mathsf{elocs}(e_{\mathrm{new}}), v \in \mathsf{Val} \end{array}$$

Fig. 6. The RDMA$^{\text{TSO}}$ program and command transitions

That is, both get and put operations result in a completion notification (cn) when complete. Indeed, this is precisely why we record an acknowledgement for each put: the ack_p serves as a placeholder for a processed put, so that we can generate an associated notification when complete.

Recall that $\mathtt{poll}(\overline{n})$ awaits the completion of the earliest unpolled get/put towards $\overline{n}$. To achieve this, completion notifications are left in $\mathbf{wb}_\mathrm{L}$ until polled. Executing $\mathtt{poll}(\overline{n})$ can proceed if the head of $\mathbf{wb}_\mathrm{L}$ of the $\overline{n}$-queue pair is a cn entry, in which case this cn entry (the earliest in FIFO order) is removed. If $\overline{n}$-$\mathbf{wb}_\mathrm{L}$ is empty or its head is a write entry, then the execution of $\mathtt{poll}(\overline{n})$ is blocked.

3.2 RDMA$^{\text{TSO}}$ Concrete Operational Semantics

We describe the RDMA$^{\text{TSO}}$ concrete operational semantics by separating the transitions of its program and hardware subsystems. The former describe the steps in program execution, e.g. how a branching program is reduced. The latter describe how the memory, store buffers and queue pairs evolve throughout the execution, e.g. how remote puts reach the memory. The RDMA$^{\text{TSO}}$ operational semantics is then defined by combining the transitions of its program and hardware subsystems.

Program Transitions. Program transitions in the middle of Fig. 6 are defined via the transitions of their constituent commands. Command transitions are of the form $\mathsf{C} \xrightarrow{l} \mathsf{C}'$, where $\mathsf{C}, \mathsf{C}' \in \mathsf{Comm}$ are sequential commands and l is a *label.* A label is either ε for silent transitions, or in Lab (defined at the top of Fig. 6) for executing a primitive command. Labels are defined as the union of Lab_n for all nodes n. A label in Lab_n may be (1) $\mathtt{lW}(x^n, v)$, for a CPU write on location x with value v; (2) $\mathtt{lR}(x^n, v)$, for a CPU read on x reading v; (3) $\mathtt{CASS}(x^n, v_1, v_2)$, for a successful CAS reading the expected value v_1 from x and updating it to v_2; (4) $\mathtt{CASF}(x^n, v)$, for a failed CAS where the expected value does not match the value v of x in memory; (5) F, for a memory fence; (6) $\mathsf{P}(\overline{n})$, for a poll on $\overline{n}$; (7) $\mathtt{Get}(x^n, y^{\overline{n}})$, for $x^n := y^{\overline{n}}$; (8) $\mathtt{Put}(y^{\overline{n}}, x^n)$, for $y^{\overline{n}} := x^n$; or (9) $\mathtt{rF}(\overline{n})$, for a remote fence on $\overline{n}$.

Command transitions for sequential composition, choice and loops (the top line) are standard. The next four transitions reduce assignments and CAS operations. As assignments and CAS involve expressions, their transitions *rewrite* expressions step-by-step, as defined at the bottom of Fig. 6 (via $\rightsquigarrow$ transitions). Each rewrite step of assignment rewrites $x := e$ as $\mathtt{assume}(y = v); x := e[v/y]$, replacing a location y in e with an arbitrary value v and checking that v matches the value of y via $\mathtt{assume}(y = v)$. Note that value v is unconstrained at this point and is later constrained when connecting program transitions with memory ones. The two rewrite transitions for CAS are analogous. Observe that when $\mathsf{C} \rightsquigarrow \mathsf{C}'$, then C silently transitions to C' (with label ε). Once e in $x := e$ is closed (i.e. contains no locations), then it is reduced to $\mathtt{skip}$ with the corresponding CPU write label for writing the value of $[\![e]\!]$ to x; *mutatis mutandis* for CAS operations. The other transitions reduce memory fences, gets, puts, remote fences, polls and assumes to $\mathtt{skip}$ with matching labels. Finally, $\mathtt{assume}(x{=}v)$ (resp. $\mathtt{assume}(x{\neq}v')$) reduces to $\mathtt{skip}$ with a $\mathtt{lR}(x, v)$ label when the in-memory value v of x matches (resp. does not match) the assertion.

Program transitions are of the form $\mathsf{P} \xrightarrow{t:l} \mathsf{P}'$, where P, P' denote (concurrent) programs, t is the identifier of the executing thread, and l is the transition label. Program transitions simply lift the transitions of their constituent threads (the bottom right rule).

Hardware Transitions. The RDMA$^{\text{TSO}}$ *hardware transitions*, given in the middle of Fig. 7, are of the form $\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:l} \mathsf{M}', \mathsf{B}', \mathsf{QP}'$, where M, M' are global *memories*, B, B' are *store-buffer maps* and $\mathsf{QP}, \mathsf{QP}'$ are *queue-pair maps*, defined at the top of Fig. 7. We model a memory as a map from locations to values. A store-buffer map is a function mapping each thread t to a *store buffer* on its associated node $n(t)$. A store buffer b on node n is a sequence of CPU writes and RDMA puts, gets and fences, as described in §3.1 (see Fig. 5). A queue-pair map is a function mapping each thread t (on node $n(t)$) to another function associating each non-$n(t)$ node with a queue pair. That is, each thread t is associated with a set of queue pairs, one for each other node on the network. A queue pair qp is a tuple of six buffers, $\mathbf{req}_\mathrm{L}$, $\mathbf{in}_\mathrm{R}$, $\mathbf{wb}_\mathrm{R}$, $\mathbf{out}_\mathrm{R}$, $\mathbf{rsp}_\mathrm{L}$ and $\mathbf{wb}_\mathrm{L}$, as described in §3.1. Each queue-pair buffer in turn is a sequence of entries as prescribed in Fig. 5. We write 'qp.' to project components of qp and use a standard map update notation to modify these components.

The first five hardware transitions describe the execution of CPU operations, as described in §3.1. Specifically, when a thread writes v to x, it records this write in its store buffer (first transition). Recall that when a thread t reads from x, it first consults its own store buffer, followed by the memory if no write to x is found in the store buffer. This lookup chain is captured by $\mathsf{M} \lhd \mathsf{B}(t)$ (second transition), defined below the hardware transitions in Fig. 7. The execution of a CAS or $\mathtt{mfence}$ proceeds if the store buffer of the executing thread is empty (the next three transitions).

The next transition describes the debuffering of a CPU write and propagating it to the memory. Similarly, the transition after describes debuffering a remote operation rc^n towards node $\overline{n}$, where it is removed from the store buffer and appended to the $\mathbf{req}_\mathrm{L}$ component of the $\overline{n}$-queue pair.

The next three transitions describe executing a remote get, put or fence operation, where a corresponding entry is added to the store buffer. The penultimate transition describes executing a poll on $\overline{n}$, which removes the earliest completion notification in $\mathbf{wb}_\mathrm{L}$ of the $\overline{n}$-queue pair. The last transition describes how the queue-pair entries travel through its six buffers, captured by the *queue-pair transitions* ($\rightarrow_{\mathsf{qp}}$) defined at the bottom of Fig. 7 (ignoring highlights for now).

Recall from §3.1 that once a get operation enters a queue pair it travels through the queue-pair pipeline in six steps, i.e. G_1–G_6 on p. 11. This is captured by the top $\rightarrow_{\mathsf{qp}}$ transition at the bottom of Fig. 7, where for brevity we have combined these six steps into one transition with a disjunctive

$$\mathsf{M} \in \mathsf{Mem} \triangleq \mathsf{Loc} \to \mathsf{Val} \qquad \mathsf{B} \in \mathsf{SBMap} \triangleq \lambda t \in \mathsf{Tid}.\mathsf{SBuff}_{n(t)} \qquad \mathsf{QP} \in \mathsf{QPMap} \triangleq \lambda t.\big(\lambda \overline{n(t)}.\mathsf{QPair}^{\overline{n}}_{n}\big)$$

$$\mathsf{b} \in \mathsf{SBuff}_n \triangleq \big\{x^n{:=}v, y^{\overline{n}}{:=}x^n, x^n{:=}y^{\overline{n}}, \mathtt{rfence}\ \overline{n}\big\}^* \qquad \mathsf{qp} \in \mathsf{QPair}^{\overline{n}}_{n} \triangleq \mathsf{Req}^{\overline{n}}_{n} \times \mathsf{In}^{\overline{n}}_{n} \times \mathsf{WBR}^{\overline{n}}_{n} \times \mathsf{Out}^{\overline{n}}_{n} \times \mathsf{Rsp}^{\overline{n}}_{n} \times \mathsf{WBL}^{\overline{n}}_{n}$$

$$\mathbf{req_L} \in \mathsf{Req}^{\overline{n}}_{n} \triangleq \big\{y^{\overline{n}}{:=}x^n, x^n{:=}y^{\overline{n}}, \mathtt{rfence}\ \overline{n}\big\}^* \quad \mathbf{in_R} \in \mathsf{In}^{\overline{n}}_{n} \triangleq \big\{y^{\overline{n}} := v, x^n := y^{\overline{n}}\big\}^* \quad \mathbf{wb_L} \in \mathsf{WBL}^{\overline{n}}_{n} \triangleq \big\{\mathsf{cn}, x^n := v\big\}^*$$

$$\mathbf{out_R} \in \mathsf{Out}^{\overline{n}}_{n} \triangleq \big\{\mathsf{ack_p}, x^n := v, x^n := y^{\overline{n}}\big\}^* \quad \mathbf{rsp_L} \in \mathsf{Rsp}^{\overline{n}}_{n} \triangleq \big\{\mathsf{ack_p}, x^n := v\big\}^* \quad \mathbf{wb_R} \in \mathsf{WBR}^{\overline{n}}_{n} \triangleq \big\{y^{\overline{n}} := v\big\}^*$$

$$\frac{\mathsf{B}' = \mathsf{B}[t \mapsto (x := v) \cdot \mathsf{B}(t)]}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\mathtt{lW}(x,v)} \mathsf{M},\mathsf{B}',\mathsf{QP}} \qquad \frac{(\mathsf{M} \lhd \mathsf{B}(t))(x) = v}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\mathtt{lR}(x,v)} \mathsf{M},\mathsf{B},\mathsf{QP}} \qquad \frac{\mathsf{B}(t) = \varepsilon \qquad \mathsf{M}(x) = v_1}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\mathtt{CASS}(x,v_1,v_2)} \mathsf{M}[x \mapsto v_2],\mathsf{B},\mathsf{QP}}$$

$$\frac{\mathsf{B}(t) = \varepsilon \qquad \mathsf{M}(x) = v}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\mathtt{CASF}(x,v)} \mathsf{M},\mathsf{B},\mathsf{QP}} \qquad \frac{\mathsf{B}(t) = \varepsilon}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\mathtt{F}} \mathsf{M},\mathsf{B},\mathsf{QP}} \qquad \frac{\mathsf{B}(t) = \mathsf{b} \cdot (x := v)}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\varepsilon} \mathsf{M}[x \mapsto v],\mathsf{B}[t \mapsto \mathsf{b}],\mathsf{QP}}$$

$$\frac{\mathsf{B}(t){=}\mathsf{b}{\cdot}\mathsf{rc}^n \qquad \mathsf{rc}^n \in \big\{x := y^{\overline{n}}, y^{\overline{n}} := x, \mathtt{rfence}\ \overline{n}\big\} \qquad \mathsf{QP}(t)(\overline{n}){=}\mathsf{qp} \qquad \mathsf{qp}' = \mathsf{qp}[\mathbf{req_L} \mapsto \mathsf{rc}^n \cdot \mathsf{qp}.\mathbf{req_L}]}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\varepsilon} \mathsf{M},\mathsf{B}[t \mapsto \mathsf{b}],\mathsf{QP}[t \mapsto \mathsf{QP}(t)[\overline{n} \mapsto \mathsf{qp}']]}$$

$$\frac{\mathsf{B}' = \mathsf{B}[t \mapsto (x := \overline{y}) \cdot \mathsf{B}(t)]}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\mathtt{Get}(x,\overline{y})} \mathsf{M},\mathsf{B}',\mathsf{QP}} \qquad \frac{\mathsf{B}' = \mathsf{B}[t \mapsto (\overline{y} := x) \cdot \mathsf{B}(t)]}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\mathtt{Put}(\overline{y},x)} \mathsf{M},\mathsf{B}',\mathsf{QP}} \qquad \frac{\mathsf{B}' = \mathsf{B}[t \mapsto (\mathtt{rfence}\ \overline{n}) \cdot \mathsf{B}(t)]}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\mathtt{rF}(\overline{n})} \mathsf{M},\mathsf{B}',\mathsf{QP}}$$

$$\frac{\mathsf{QP}(t)(\overline{n}){=}\mathsf{qp} \qquad \mathsf{qp}.\mathbf{wb_L}{=}\alpha \cdot \mathsf{cn} \qquad \mathsf{qp}'{=}\mathsf{qp}[\mathbf{wb_L} \mapsto \alpha]}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\mathtt{P}(\overline{n})} \mathsf{M},\mathsf{B},\mathsf{QP}[t \mapsto \mathsf{QP}(t)[\overline{n} \mapsto \mathsf{qp}']]} \qquad \frac{\mathsf{M},\mathsf{QP}(t)(\overline{n}) \to_{\mathsf{qp}} \mathsf{M}',\mathsf{qp}}{\mathsf{M},\mathsf{B},\mathsf{QP} \xrightarrow{t:\varepsilon} \mathsf{M}',\mathsf{B},\mathsf{QP}[t \mapsto \mathsf{QP}(t)[\overline{n} \mapsto \mathsf{qp}]]}$$

$$\text{with} \quad (\mathsf{M} \lhd \alpha)(x) \triangleq \begin{cases} v & \text{if } \alpha = \beta \cdot (x := v) \cdot - \land \forall v'.\, x := v' \notin \beta \\ \mathsf{M}(x) & \text{if } \forall v.\, x := v \notin \alpha \end{cases}$$

$$\frac{\begin{array}{llll} & \mathsf{qp}.\mathbf{req_L} = \alpha \cdot (x{:=}\overline{y}) & \mathsf{qp}' = \mathsf{qp}[\mathbf{req_L} \mapsto \alpha][\mathbf{in_R} \mapsto (x{:=}\overline{y}) \cdot \mathsf{qp}.\mathbf{in_R}] & \mathsf{M}' = \mathsf{M} \\ \lor & \mathsf{qp}.\mathbf{in_R} = \alpha \cdot (x{:=}\overline{y}) & \mathsf{qp}' = \mathsf{qp}[\mathbf{in_R} \mapsto \alpha][\mathbf{out_R} \mapsto (x{:=}\overline{y}) \cdot \mathsf{qp}.\mathbf{out_R}] & \mathsf{M}' = \mathsf{M} \\ \lor & \mathsf{qp}.\mathbf{out_R} = \alpha \cdot (x{:=}\overline{y}) \cdot \beta & \mathsf{qp}.\mathbf{wb_R} = \varepsilon \qquad \mathsf{qp}' = \mathsf{qp}[\mathbf{out_R} \mapsto \alpha \cdot (x{:=}\mathsf{M}(\overline{y})) \cdot \beta] & \mathsf{M}' = \mathsf{M} \\ \lor & \mathsf{qp}.\mathbf{out_R} = \alpha \cdot (x{:=}v) & \mathsf{qp}' = \mathsf{qp}[\mathbf{out_R} \mapsto \alpha][\mathbf{rsp_L} \mapsto (x{:=}v) \cdot \mathsf{qp}.\mathbf{rsp_L}] & \mathsf{M}' = \mathsf{M} \\ \lor & \mathsf{qp}.\mathbf{rsp_L} = \alpha \cdot (x{:=}v) & \mathsf{qp}' = \mathsf{qp}[\mathbf{rsp_L} \mapsto \alpha][\mathbf{wb_L} \mapsto \mathsf{cn} \cdot (x{:=}v) \cdot \mathsf{qp}.\mathbf{wb_L}] & \mathsf{M}' = \mathsf{M} \\ \lor & \mathsf{qp}.\mathbf{wb_L} = \alpha \cdot (x{:=}v) \cdot \mathsf{cn}^* & \mathsf{qp}' = \mathsf{qp}[\mathbf{wb_L} \mapsto \alpha \cdot \mathsf{cn}^*] & \mathsf{M}' = \mathsf{M}[x \mapsto v] \end{array}}{\mathsf{M},\mathsf{qp} \to_{\mathsf{qp}} \mathsf{M}',\mathsf{qp}'}$$

$$\frac{\begin{array}{llll} & \mathsf{qp}.\mathbf{req_L} = \alpha \cdot (\overline{y}{:=}x) & \mathsf{qp}.\mathbf{wb_L} = \mathsf{cn}^* \qquad \mathsf{qp}' = \mathsf{qp}[\mathbf{req_L} \mapsto \alpha][\mathbf{in_R} \mapsto (\overline{y}{:=}\,\mathsf{M}(x)) \cdot \mathsf{qp}.\mathbf{in_R}] & \mathsf{M}' = \mathsf{M} \\ \lor & \mathsf{qp}.\mathbf{in_R} = \alpha \cdot (\overline{y}{:=}v) & \mathsf{qp}' = \mathsf{qp}[\mathbf{in_R} \mapsto \alpha][\mathbf{wb_R} \mapsto (\overline{y}{:=}v) \cdot \mathsf{qp}.\mathbf{wb_R}][\mathbf{out_R} \mapsto \mathsf{ack_p} \cdot \mathsf{qp}.\mathbf{out_R}] & \mathsf{M}' = \mathsf{M} \\ \lor & \mathsf{qp}.\mathbf{wb_R} = \alpha \cdot (\overline{y}{:=}v) & \mathsf{qp}' = \mathsf{qp}[\mathbf{wb_R} \mapsto \alpha] & \mathsf{M}'{=}\mathsf{M}[\overline{y} \mapsto v] \\ \lor & \mathsf{qp}.\mathbf{out_R} = \alpha \cdot \mathsf{ack_p} & \mathsf{qp}' = \mathsf{qp}[\mathbf{out_R} \mapsto \alpha][\mathbf{rsp_L} \mapsto \mathsf{ack_p} \cdot \mathsf{qp}.\mathbf{rsp_L}] & \mathsf{M}' = \mathsf{M} \\ \lor & \mathsf{qp}.\mathbf{rsp_L} = \alpha \cdot \mathsf{ack_p} & \mathsf{qp}' = \mathsf{qp}[\mathbf{rsp_L} \mapsto \alpha][\mathbf{wb_L} \mapsto \mathsf{cn} \cdot \mathsf{qp}.\mathbf{wb_L}] & \mathsf{M}' = \mathsf{M} \end{array}}{\mathsf{M},\mathsf{qp} \to_{\mathsf{qp}} \mathsf{M}',\mathsf{qp}'}$$

$$\frac{\mathsf{qp}.\mathbf{req_L}{=}\alpha \cdot (\mathtt{rfence}\ \overline{n}) \qquad \mathsf{qp}.\mathbf{in_R} = \mathsf{qp}.\mathbf{out_R} = \mathsf{qp}.\mathbf{rsp_L} = \varepsilon \qquad \mathsf{qp}' = \mathsf{qp}[\mathbf{req_L} \mapsto \alpha]}{\mathsf{M},\mathsf{qp} \to_{\mathsf{qp}} \mathsf{M},\mathsf{qp}'}$$

Fig. 7. RDMA$^{\mathsf{TSO}}$ hardware domains (above), hardware transitions (middle) and queue-pair transitions (below)

premise, with each disjunct corresponding to a step in G$_1$–G$_6$. That is, if any of the six disjuncts in the premise of the rule hold, then the transition is enabled. Analogously, the second $\to_{\mathsf{qp}}$ transition describes how a put operation proceeds through the queue-pair pipeline, with each of the five disjuncts describing one of the five steps in P$_1$–P$_5$ (p. 11). Finally, the last $\to_{\mathsf{qp}}$ transition describes

$$\frac{\mathsf{P} \xrightarrow{t:\varepsilon} \mathsf{P}'}{\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP} \Rightarrow \mathsf{P}', \mathsf{M}, \mathsf{B}, \mathsf{QP}} \qquad \frac{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\varepsilon} \mathsf{M}', \mathsf{B}', \mathsf{QP}'}{\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP} \Rightarrow \mathsf{P}, \mathsf{M}', \mathsf{B}', \mathsf{QP}'} \qquad \frac{\mathsf{P} \xrightarrow{t:l} \mathsf{P}' \quad \mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:l} \mathsf{M}', \mathsf{B}', \mathsf{QP}'}{\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP} \Rightarrow \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}'}$$

Fig. 8. RDMA$^{\text{TSO}}$ operational semantics with the program and hardware transitions given in Fig. 6 and Fig. 7

the execution of a remote fence as described in §3.1, ensuring that it can only proceed if there are no pending operations on the queue pair (i.e. $\mathsf{qp}.\mathbf{in_R} = \mathsf{qp}.\mathbf{out_R} = \mathsf{qp}.\mathbf{rsp_L} = \varepsilon$).

RDMA$^{\text{TSO}}$ Combined Transitions. The RDMA$^{\text{TSO}}$ operational semantics is defined by combining its program and hardware transitions as shown in Fig. 8. When the program subsystem takes a silent step, then the hardware subsystem is unchanged (first transition); analogously, when the hardware subsystem takes a silent step, then the program subsystem is unchanged (second transition). Finally, when the program and hardware subsystems both take the same transition (with the same label and by the same thread), then the transition effect is that of their combined effects.

Operational Semantics without PCIe. Recall from §2 (p. 8) that we model the PCIe-specific guarantee where a NIC remote read propagates all pending NIC remote writes (in $\mathbf{wb_R}$) to memory, while a NIC local read propagates all pending NIC local writes (in $\mathbf{wb_L}$) to memory. Nevertheless, we can relax this as follows. For NIC remote reads, we can replace the highlighted premises of the get queue-pair transition (top $\longrightarrow_{\mathsf{qp}}$ transition) in Fig. 7 with $\mathsf{qp}' = \mathsf{qp}[\mathbf{out_R} \mapsto \alpha \cdot (x{:=}(\mathsf{M} \lhd \mathbf{wb_R})(\overline{y})) \cdot \beta]$. That is, we no longer require $\mathbf{wb_R}$ to be empty (i.e. there may be pending writes in $\mathbf{wb_R}$), and when reading the value of $\overline{y}$ we first check for pending writes on $\overline{y}$ in $\mathbf{wb_R}$.

For NIC local reads, we can similarly replace the highlighted premises of the put queue-pair transition (middle $\longrightarrow_{\mathsf{qp}}$ transition) with $\mathsf{qp}' = \mathsf{qp}[\mathbf{req_L} \mapsto \alpha][\mathbf{in_R} \mapsto (\overline{y} := (\mathsf{M} \lhd \mathbf{wb_L})(x)) \cdot \mathsf{qp}.\mathbf{in_R}]$. That is, we no longer require $\mathbf{wb_L}$ to contain only completion notifications (and allow it also to contain pending writes), and we first check $\mathbf{wb_L}$ for pending writes when reading the value of x.

Observations. Given a thread t and its store buffer b, all remote operations by t also go through b. Hence, as store buffers are FIFO, a CPU write cc in b before a remote operation rc in b always reaches the memory before rc is debuffered, and thus cc is visible to rc. Moreover, as all six buffers of the queue-pair pipeline are FIFO, remote operations maintain the order in which they were issued as they go through the queue-pair pipeline. Therefore, a thread always receives the completion notifications of get/put operations in the order they were submitted (i.e. the program order).

Observe that an `rFence` stipulates that only $\mathbf{in_R}$, $\mathbf{out_R}$ and $\mathbf{rsp_L}$ be empty but not $\mathbf{wb_L}$ and $\mathbf{wb_R}$. As such, `rFence` cannot guarantee that the result of earlier put (resp. get) operations have reached the remote (resp. local) memory: they can still be pending in $\mathbf{wb_R}$ (resp. $\mathbf{wb_L}$). Moreover, `rFence` has no bearing on CPU operations and does not block their execution. Hence, later CPU operations (after `rFence`) may be visible to earlier get/put operations (those before `rFence`).

Recall that a put operation $x^n := y$ comprises a local read from y and a remote write to x, and a get operation $x := y^n$ comprises a remote read from y and a local write to x. Note that the local read of a put is fulfilled when it reaches the head of $\mathbf{in_R}$ and is subsequently forwarded to $\mathbf{out_R}$. This ensures that puts are executed in program order and that puts are executed before all later gets (as in Fig. 2g). By contrast, the remote read of a get is fulfilled while in $\mathbf{out_R}$ *non-deterministically* (i.e. not necessarily when it is at the head of $\mathbf{out_R}$). This means that remote reads of gets can be *reordered* with respect to one another, as well as with respect to the remote writes of puts (as in Fig. 2h). Such reorderings can be prevented by adding an `rFence` after a get (as in Fig. 2i).

Lastly, note that a `poll` retrieves the earliest `cn` in $\mathbf{wb}_\mathrm{L}$ (i.e. at its head). In the case of gets, the result of the get (its local write) is sent to $\mathbf{wb}_\mathrm{L}$ before its associated `cn`. As such, if the head of $\mathbf{wb}_\mathrm{L}$ is `cn`, then its result is guaranteed to have reached the memory of the local node when polling. By contrast, in the case of puts, its remote write operation could still be in $\mathbf{wb}_\mathrm{R}$ when polling, and thus polling a put cannot guarantee that the effect of the put has reached the remote memory.

3.3 RDMA$^{\text{TSO}}$ Simplified Operational Semantics

The concrete operational semantics in §3.2 reflects the structure of the underlying hardware, namely that of the six buffers in a queue pair. However, since all four buffers ($\mathbf{req}_\mathrm{L}$, $\mathbf{in}_\mathrm{R}$, $\mathbf{out}_\mathrm{R}$, $\mathbf{rsp}_\mathrm{L}$) in the middle are FIFO, we can simplify them by modelling them as a *single* buffer. Specifically, we model a *simple queue pair*, $\mathbf{sqp} \in \mathsf{SQPair}^{\overline{n}}_n \triangleq \mathsf{Pipe}^{\overline{n}}_n \times \mathsf{WBR}^{\overline{n}}_n \times \mathsf{WBL}^{\overline{n}}_n$, as a tuple comprising a *pipe*, as well as local and remote write-back buffers (as before). A pipe, $\mathbf{pipe} \in \mathsf{Pipe}^{\overline{n}}_n \triangleq \left\{y^{\overline{n}} := x^n, y^{\overline{n}} := v, \mathsf{ack}_\mathrm{p}, x^n := y^{\overline{n}}, x^n := v, \mathtt{rfence}\ \overline{n}\right\}^*$, is simply a sequence of puts, gets, simplified puts and gets (with values), acknowledgements, and remote fences.

The program and command transitions of our simplified semantics is identical to those of the concrete one (Fig. 6). Moreover, the hardware transitions of our simplified semantics are those of the concrete semantics (Fig. 7), except that the queue-pair transitions at the bottom of Fig. 7 are replaced with the *simplified queue-pair transitions* given in §B. As before, the simplified operational semantics is obtained by combining its program and hardware transitions (Fig. 8).

Finally, we show that our concrete operational semantics is equivalent to the simplified one. We have mechanised our proof in Coq, available in the supplementary material [Ambal et al. 2024], with an overview of the proof available in §B.

Theorem 3.1. *The concrete RDMA$^{\text{TSO}}$ operational semantics is equivalent to the simplified one.*

4 RDMA$^{\text{TSO}}$ DECLARATIVE SEMANTICS

Events and Executions. In the literature of declarative models, the traces of a program are commonly represented as a set of *executions*, where an execution is a graph comprising: i) a set of *events* (graph nodes); and ii) a number of relations on events (graph edges). Each event is associated with the execution of a primitive command (in PComm) and is a tuple (ι, t, l), where ι is the (unique) *event identifier*, $t \in \mathsf{Tid}$ identifies the executing thread, and $l \in \mathsf{ELab}$ is the *event label*, defined below.

Definition 4.1 (Labels and events). An *event*, $\mathsf{e} \in \mathsf{Event}$, is a triple (ι, t, l), where $\iota \in \mathbb{N}$, $t \in \mathsf{Tid}$ and $l \in \mathsf{ELab}_{n(t)}$. The set of *event labels* is $\mathsf{ELab} \triangleq \bigcup_n \mathsf{ELab}_n$ for all nodes n. An *event label* of n, $l \in \mathsf{ELab}_n$, is a tuple of one of the following forms:

- (CPU) local read: $l = \mathtt{lR}(x^n, v_\mathrm{r})$
- (CPU) local write: $l = \mathtt{lW}(x^n, v_\mathrm{w})$
- (CPU) CAS: $l = \mathtt{CAS}(x^n, v_\mathrm{r}, v_\mathrm{w})$
- (CPU) memory fence: $l = \mathtt{F}$
- (CPU) poll: $l = \mathtt{P}(\overline{n})$
- NIC local read: $l = \mathtt{nlR}(x^n, v_\mathrm{r}, \overline{n})$
- NIC remote write: $l = \mathtt{nrW}(y^{\overline{n}}, v_\mathrm{w})$
- NIC remote read: $l = \mathtt{nrR}(y^{\overline{n}}, v_\mathrm{r})$
- NIC local write: $l = \mathtt{nlW}(x^n, v_\mathrm{w}, \overline{n})$
- NIC fence: $l = \mathtt{nF}(\overline{n})$

Each event label denotes whether the associated primitive command is handled by the NIC (right column, prefixed with n), or the CPU (left column). A poll instruction is handled by the CPU (it simply awaits for a completion notification from the NIC). A put operation $x^n := y$, which consists of a NIC local read from y and a NIC remote write to x, is modelled as two events of type $\mathtt{nlR}$ (on y) and $\mathtt{nrW}$ (on x). Similarly, a get $x := y^n$ is modelled as two events of type $\mathtt{nrR}$ (on y) and $\mathtt{nlW}$ (on x).

We write $\mathtt{type}(l)$, $\mathtt{loc}(l)$, $v_\mathrm{r}(l)$, $v_\mathrm{w}(l)$, $n(l)$ and $\overline{n}(l)$ for the type (e.g. $\mathtt{lR}$), location, read value, write value, node and remote node of l, where applicable; e.g. $\mathtt{loc}(\mathtt{nlR}(x^n, v_\mathrm{r}, \overline{n})) = x^n$, $n(\mathtt{nlR}(x^n, v_\mathrm{r}, \overline{n})) = n$

and $\overline{n}(\mathtt{nlR}(x^n, v_\mathrm{r}, \overline{n})) = \overline{n}$. We lift these functions to events as expected. We write $\iota(\mathsf{e})$, $t(\mathsf{e})$, $l(\mathsf{e})$ to project the corresponding components of an event $\mathsf{e} = (\iota, t, l)$.

Issue and Observation Points. In what follows we distinguish between when an instruction is *issued* and when it is *observed.* Intuitively, an instruction is issued when it is processed by the CPU or the NIC, and it is observed when its effect is propagated to the local or remote memory. As such, since writes (be they by the CPU or NIC) are the only instructions that have an observable effect on memory, the time points at which they are issued and observed may differ. A CPU write is issued when it is added to the store buffer and is observed when it is debuffered and propagated to memory. Similarly, the local (resp. remote) write of a get (resp. put) is issued when it is added to $\mathbf{wb_L}$ (resp. $\mathbf{wb_R}$), and observed when it is propagated to memory. By contrast, all other events are *instantaneous* in that *either* they do not have an observable effect on memory and thus their issue and observation points coincide, *or* their effect is written to memory *immediately.* In particular, CAS operations are instantaneous. Note that the observation point of any instruction either coincides with its issue point (instantaneous events) or it follows its issue point (write events).

Notation. Given a relation r and a set A, we write r^+ for the transitive closure of r; r^{-1} for the inverse of r; $r|_A$ for $r \cap (A \times A)$; and $[A]$ for the identity relation on A, i.e. $\{(a, a) \mid a \in A\}$. We write $r_1; r_2$ for the relational composition of r_1 and r_2: $\{(a, b) \mid \exists c.\ (a, c) \in r_1 \land (c, b) \in r_2\}$. When r is a strict partial order, we write $r|_\mathrm{imm}$ for the *immediate* edges in r, i.e. $r \setminus (r; r)$. Given a set of events E and a location x, we write E_x for $\{\mathsf{e} \in E \mid \mathtt{loc}(\mathsf{e}) = x\}$. Given a set of events E and a label type $\mathtt{X}$, we write $E.\mathtt{X}$ for $\{\mathsf{e} \in E \mid \mathsf{type}(\mathsf{e}) = \mathtt{X}\}$, and define its sets of *reads* as $E.\mathcal{R} \triangleq E.\mathtt{lR} \cup E.\mathtt{CAS} \cup E.\mathtt{nlR} \cup E.\mathtt{nrR}$, *writes* as $E.\mathcal{W} \triangleq E.\mathtt{lW} \cup E.\mathtt{CAS} \cup E.\mathtt{nlW} \cup E.\mathtt{nrW}$, *NIC writes* as $E.\mathtt{nW} \triangleq E.\mathtt{nlW} \cup E.\mathtt{nrW}$ and *instantaneous events* as $E.\mathtt{Inst} \triangleq E \setminus (E.\mathtt{lW} \cup E.\mathtt{nlW} \cup E.\mathtt{nrW})$. Intuitively, the effects of CPU writes, NIC local writes and NIC remote writes (labelled $\mathtt{lW}$, $\mathtt{nlW}$ and $\mathtt{nrW}$) are only visible when they respectively leave the store buffer, $\mathbf{wb_L}$, and $\mathbf{wb_R}$, and are thus excluded from the set of instantaneous events.

The ‘*same-location*’ relation is $\mathsf{sloc} \triangleq \{(\mathsf{e}, \mathsf{e}') \in \mathsf{Event}^2 \mid \mathtt{loc}(\mathsf{e}) = \mathtt{loc}(\mathsf{e}')\}$; the ‘*same-thread*’ relation is $\mathsf{sthd} \triangleq \{(\mathsf{e}, \mathsf{e}') \in \mathsf{Event}^2 \mid t(\mathsf{e}) = t(\mathsf{e}')\}$; and the ‘*same-queue-pair*’ relation is $\mathsf{sqp} \triangleq \{(\mathsf{e}, \mathsf{e}') \in \mathsf{Event}^2 \mid t(\mathsf{e}) = t(\mathsf{e}') \land \overline{n}(\mathsf{e}) = \overline{n}(\mathsf{e}')\}$. Note that $\mathsf{sqp} \subseteq \mathsf{sthd}$ and that sloc, sthd and sqp are all symmetric. For a set of events E, we write $E.\mathsf{sloc}$ for $\mathsf{sloc}|_E$; similarly for $E.\mathsf{sthd}$ and $E.\mathsf{sqp}$.

Definition 4.2 (Pre-executions). A *pre-execution* is a tuple $G = \langle E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathsf{pf}, \mathsf{nfo} \rangle$, where:

- $E \subseteq \mathsf{Event}$ is the set of events and includes a set of *initialisation* events, $E^0 \subseteq E$, comprising a single write with label $\mathtt{lW}(x, 0)$ for each $x \in \mathsf{Loc}$.
- $\mathsf{po} \subseteq E \times E$ is the ‘*program order*’ relation defined as a disjoint union of strict total orders, each ordering the events of one thread, with $E^0 \times (E \setminus E^0) \subseteq \mathsf{po}$.
- $\mathsf{rf} \subseteq E.\mathcal{W} \times E.\mathcal{R}$ is the ‘*reads-from*’ relation on events of the same location with matching values; i.e. $(a, b) \in \mathsf{rf} \Rightarrow (a, b) \in \mathsf{sloc} \land v_\mathrm{w}(a) = v_\mathrm{r}(b)$. Moreover, rf is total and functional on its range: every read in $E.\mathcal{R}$ is related to exactly one write in $E.\mathcal{W}$.
- $\mathsf{mo} \triangleq \bigcup_{x \in \mathsf{Loc}} \mathsf{mo}_x$ is the ‘*modification-order*’, where each mo_x is a strict total order on $E.\mathcal{W}_x$ with $E^0_x \times (E.\mathcal{W}_x \setminus E^0_x) \subseteq \mathsf{mo}_x$ describing the order in which writes on x reach the memory.
- $\mathsf{pf} \subseteq E.\mathtt{nW} \times E.\mathtt{P}$ is the ‘*polls-from*’ relation, relating earlier (in program-order) NIC writes to later poll operations on the *same queue pair*; i.e. $\mathsf{pf} \subseteq \mathsf{po} \cap \mathsf{sqp}$. Moreover, pf is functional on its domain (every NIC write can be be polled at most once), and pf is total and functional on its range (every poll in $E.\mathtt{P}$ polls from exactly one NIC write).
- $\mathsf{nfo} \subseteq E.\mathsf{sqp}$ is the ‘*NIC flush order*’, such that for all $(a, b) \in E.\mathsf{sqp}$, if $a \in E.\mathtt{nlR}$, $b \in E.\mathtt{nlW}$, then $(a, b) \in \mathsf{nfo} \cup \mathsf{nfo}^{-1}$, and if $a \in E.\mathtt{nrR}$, $b \in E.\mathtt{nrW}$, then $(a, b) \in \mathsf{nfo} \cup \mathsf{nfo}^{-1}$.

Recall from §2 that we model the PCIe-specific guarantee where a NIC local read ($\mathtt{nlR}$) propagates all pending NIC local writes ($\mathtt{nlW}$) in $\mathbf{wb_L}$ (on the same queue pair) to memory, while a NIC remote read (nrR) propagates all pending NIC remote writes (nrW) in $\mathbf{wb_R}$ (on the same queue pair) to memory. In other words, the issue point of an nlR event e_r is totally ordered with respect to the issue and observation points of any nlW event e_w on the same queue pair. Specifically, either (1) e_w had already been flushed before issuing e_r (i.e. e_w had already left $\mathbf{wb_L}$) and thus e_w is issued and observed before e_r; or (2) e_w was in $\mathbf{wb_L}$ before issuing e_r, in which case we cannot issue e_r until $\mathbf{wb_L}$ is emptied, propagating e_w to memory, and thus e_w is issued and observed before e_r; or (3) e_w has not reached $\mathbf{wb_L}$ when we issue e_r, in which case e_w will reach and leave $\mathbf{wb_L}$ (i.e. is issued and observed) after e_r. Similarly for nrR/nrW events. We model this total order through the nfo relation, stipulating that all NIC local reads and writes (resp. all NIC remote reads and writes) on the same queue pair be totally ordered. Later we describe how we can relax this PCIe guarantee (see p. 21).

Derived Relations. Given a pre-execution $\langle E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathsf{pf}, \mathsf{nfo}\rangle$, we define the '*reads-before*' relation as $\mathsf{rb} \triangleq (\mathsf{rf}^{-1}; \mathsf{mo}) \setminus [E]$, relating each read r to writes that are mo-after the write r reads from. We further define the rf-*buffer* relation as $\mathsf{rf_b} \triangleq [\mathtt{lW}]; (\mathsf{rf} \cap \mathsf{sthd}); [\mathtt{lR}]$, including CPU rf edges by the same thread (with access to the same store buffer). We define the $\mathsf{rf_b}$-complement as $\mathsf{rf}_{\overline{\mathsf{b}}} \triangleq \mathsf{rf} \setminus \mathsf{rf_b}$, including all other rf edges (i.e. by different threads or involving remote operations). Intuitively, when $w \xrightarrow{\mathsf{rf_b}} r$, then r may read from w *before it is observable*. Specifically, as CPU writes are delayed in the store buffer and CPU reads first check the buffer, r can read from w either when (1) w is still in the thread's store buffer (i.e. w is not yet observable); or (2) w is in the memory (i.e. w is observable). By contrast, when $w \xrightarrow{\mathsf{rf}_{\overline{\mathsf{b}}}} r$, then r reads from w only once it is observable (i.e. it has reached the memory). Analogously, we define the rb-*buffer* relation as $\mathsf{rb_b} \triangleq [\mathtt{lR}]; (\mathsf{rb} \cap \mathsf{sthd}); [\mathtt{lW}]$.

Recall that we distinguish between the issue and observation points of an event. We thus define the '*issue-preserved program order*', ippo, as the subset of po edges ($\mathsf{ippo} \subseteq \mathsf{po}$) that must be preserved when issuing instructions. That is, if two events are ippo-related, then they must be issued in program order; otherwise they may be reordered and thus issued in either order. The table at the top of Fig. 9 describes which po edges are included in ippo, where ✓ denotes that the two instructions are ippo-related (i.e. their issue order is preserved and they must be issued in program order), ✗ denotes that they are not ippo-related (i.e. their issue order is not preserved and they may be issued out of order) and sqp denotes that they are ippo-related iff they are on the same queue pair. For instance, when a CPU instruction is followed by a NIC one, then they are issued in order (top-right quadrant of the table). By contrast, when a NIC instruction is followed by a CPU one, then they may be reordered (bottom-left quadrant), *as if* the NIC instruction was executed in a parallel thread (as discussed in §2), resulting in weak behaviours such as $z=1$ in Fig. 2b.

Analogously, we define the '*observation-preserved program order*', oppo, as the subset of po edges ($\mathsf{oppo} \subseteq \mathsf{po}$) that must be preserved when observing the effects of instructions. In other words, if two events are oppo-related, then they become observable in program order; otherwise they may be reordered and become observable in either order. The table at the bottom of Fig. 9 describes which po edges are included in oppo, with ✓, ✗ and sqp interpreted in the same way as for ippo.

Observe that ippo and oppo only differ in four cells. In the case of B1 and B5, this is because CPU writes (type lW) are delayed in the store buffer and may be reordered and thus become observable after CPU operations that do not go through the store buffer, namely CPU reads and polls. Analogously, in the case of G10 (resp. I10), this is because NIC remote (resp. local) writes are delayed in $\mathbf{wb_R}$ (resp. $\mathbf{wb_L}$), so a later remote fence do not ensure the writes have propagated to memory.

From Programs to Executions. The *semantics* of a program P is a set of *consistent* executions (defined shortly) of P, defined by induction on the structure of P. This definition is standard and omitted (see §C). The executions produced by this construction are *well-formed*, as we define below.

Later in Program Order

ippo				CPU					NIC				
				1	2	3	4	5	6	7	8	9	10
				lR	lW	CAS	F	P	nlR	nrW	nrR	nlW	nF
Earlier in Program Order	CPU	A	lR	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
		B	lW	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
		C	CAS	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
		D	F	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
		E	P	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	NIC	F	nlR	✗	✗	✗	✗	✗	sqp	sqp	sqp	sqp	sqp
		G	nrW	✗	✗	✗	✗	✗	✗	sqp	sqp	sqp	sqp
		H	nrR	✗	✗	✗	✗	✗	✗	✗	✗	sqp	sqp
		I	nlW	✗	✗	✗	✗	✗	✗	✗	✗	sqp	sqp
		J	nF	✗	✗	✗	✗	✗	sqp	sqp	sqp	sqp	sqp

Later in Program Order

oppo				CPU					NIC				
				1	2	3	4	5	6	7	8	9	10
				lR	lW	CAS	F	P	nlR	nrW	nrR	nlW	nF
Earlier in Program Order	CPU	A	lR	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
		B	lW	✗	✓	✓	✓	✗	✓	✓	✓	✓	✓
		C	CAS	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
		D	F	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
		E	P	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	NIC	F	nlR	✗	✗	✗	✗	✗	sqp	sqp	sqp	sqp	sqp
		G	nrW	✗	✗	✗	✗	✗	✗	sqp	sqp	sqp	✗
		H	nrR	✗	✗	✗	✗	✗	✗	✗	✗	sqp	sqp
		I	nlW	✗	✗	✗	✗	✗	✗	✗	✗	sqp	✗
		J	nF	✗	✗	✗	✗	✗	sqp	sqp	sqp	sqp	sqp

Fig. 9. The RDMA$^{\text{TSO}}$ ordering constraints on ippo (above) and oppo (below), where ✓ denotes that instructions are ordered (and cannot be reordered), ✗ denotes they are not ordered (and may be reordered), and sqp denotes they are ordered iff they are on the same queue pair. Replacing sqp in the highlighted cells (G8-G9) with ✗ would yield oppo for the weaker model without the PCIe guarantee of reads flushing buffers (p. 21).

Definition 4.3 (Executions). A pre-execution $G = \langle E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathsf{pf}, \mathsf{nfo}\rangle$ is *well-formed* if the following hold for all w, r, w_1, w_2, p_2:

(1) Poll events poll-from the oldest non-polled remote operation on the same queue pair: if $w_1 \in G.\mathtt{nW}$ and $w_1 \xrightarrow{\mathsf{po}\cap\mathsf{sqp}} w_2 \xrightarrow{\mathsf{pf}} p_2$, then there exists p_1 such that $w_1 \xrightarrow{\mathsf{pf}} p_1 \xrightarrow{\mathsf{po}} p_2$.

(2) Each put (resp. get) operation corresponds to two events: a read and a write with the read immediately preceding the write in po: (a) if $r \in G.\mathtt{nlR}$ (resp. $r \in G.\mathtt{nrR}$), then $(r, w) \in \mathsf{po}|_{\mathrm{imm}}$ for some $w \in G.\mathtt{nrW}$ ($w \in G.\mathtt{nlW}$); and (b) if $w \in G.\mathtt{nrW}$ (resp. $w \in G.\mathtt{nlW}$), then $(r, w) \in \mathsf{po}|_{\mathrm{imm}}$ for some $r \in G.\mathtt{nlR}$ ($r \in G.\mathtt{nrR}$).

(3) Read and write events of a put (resp. get) have matching values: if $(r, w) \in G.\mathsf{po}|_{\mathrm{imm}}$, $\mathrm{type}(r) \in \{\mathtt{nlR}, \mathtt{nrR}\}$ and $\mathrm{type}(w) \in \{\mathtt{nlW}, \mathtt{nrW}\}$, then $v_{\mathrm{r}}(r) = v_{\mathrm{w}}(w)$.

An *execution* is a pre-execution (Def. 4.2) that is well-formed.

We use 'G.' to project the components and (derived) relations of execution G; e.g. G.rf and G.ippo. When the choice of G is clear, we simply write e.g. rf and ippo. See §C.1 for execution examples.

Consistency. Note that the notion of an execution (Def. 4.3) imposes very few constraints on its po, rf, mo, pf and nfo relations. Such restrictions and thus the permitted behaviours of a program are determined by defining the set of *consistent* executions, defined below.

Definition 4.4 (RDMA^{TSO}-consistency). An execution $\langle E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathsf{pf}, \mathsf{nfo}\rangle$ is *RDMA^{TSO}-consistent* iff (1) ib is irreflexive; (2) ob is irreflexive; and (3) $([\mathtt{Inst}]; \mathsf{ib}; \mathsf{ob})^+$ is irreflexive, where:

$$\mathsf{ib} \triangleq \left(\mathsf{ippo} \cup \mathsf{rf} \cup \mathsf{pf} \cup \mathsf{nfo} \cup \mathsf{rb}_\mathsf{b}\right)^+ \qquad \text{(‘issued-before’)}$$

$$\mathsf{ob} \triangleq \left(\mathsf{oppo} \cup \mathsf{rf}_{\overline{\mathsf{b}}} \cup ([\mathtt{nlW}]; \mathsf{pf}) \cup \mathsf{nfo} \cup \mathsf{rb} \cup \mathsf{mo}\right)^+ \qquad \text{(‘observed-before’)}$$

The ib (resp. ob) relation is an extension of ippo (resp. oppo), describing the issue (resp. observation) order across the instructions of different threads and nodes. RDMA^{TSO}-consistency requires that ib and ob be irreflexive (i.e. yield strict partial orders as they are defined transitively).

The rf (resp. pf) component in ib states that if e reads from (resp. polls from) w, then w must have been issued before e. Recall that nfo totally orders the `nlR`/`nlW` and `nrR`/`nrW` operations on the same queue pair and is thus in ib. The rb_b component in ib ensures that a CPU read r by thread t on location x observes all earlier writes on x by t: if r in t reads from w_r ($w_r \xrightarrow{\mathsf{rf}} r$), which is later overwritten by w in t ($w_r \xrightarrow{\mathsf{mo}} w$ and $(w, r) \in \mathsf{sthd}$), i.e. $r \xrightarrow{\mathsf{rb}_\mathsf{b}} w$, then w must have been issued after r, as otherwise r should have read from the later w and not w_r. Note that this is not the case for $\mathsf{rb} \setminus \mathsf{rb}_\mathsf{b}$: r can be issued *after* a later write w and still not observe it because the effect of w has not yet reached the memory (w is in $\mathbf{wb}_\mathbf{L}$/$\mathbf{wb}_\mathbf{R}$ or the store buffer of another thread).

The $\mathsf{rf}_{\overline{\mathsf{b}}}$ component in ob states that if a read r reads from a write w that passed through a different buffer (i.e. either w is part of an RDMA operation and went through $\mathbf{wb}_\mathbf{L}$/$\mathbf{wb}_\mathbf{R}$, or w is a CPU write by another thread and thus went through a different store buffer), then r can only read from w once it has reached the memory, i.e. only once w is observable, and thus w must have become observable before r. The $[\mathtt{nlW}]; \mathsf{pf}$ component states that if p polls from a NIC local write w, then w must have left the $\mathbf{wb}_\mathbf{L}$ buffer and reached the memory. Note that this is not the case for `nrW` events: polling an `nrW` event w succeeds when w is in $\mathbf{wb}_\mathbf{R}$ and w may not have yet reached the memory. The nfo in ob can be justified as in the case of ib.

The rb component ensures that a read r on location x observes the latest write on x that has reached the memory: if $w_r \xrightarrow{\mathsf{rf}} r$ and $w_r \xrightarrow{\mathsf{mo}} w$, i.e. $r \xrightarrow{\mathsf{rb}} w$, then w must have reached the memory after r was issued/observed, as otherwise r should have read from this later w and not w_r. As mo describes the order in which the writes on each location reach the memory, it is included in ob.

The third condition Def. 4.4 asks that $([\mathtt{Inst}]; \mathsf{ib}; \mathsf{ob})^+$ be irreflexive. Intuitively, if $\mathsf{e}_1 \xrightarrow{\mathsf{ib}} \mathsf{e}_2 \xrightarrow{\mathsf{ob}} \mathsf{e}_3$, then (1) e_1 is issued before e_2; (2) e_2 is issued before it is observed; and (3) e_2 is observed before e_3. If e_1 and e_3 are instantaneous events (i.e. their issue and observation points coincide) then e_1 necessarily precedes e_3, and thus the semantics prohibits creating a cycle from these dependencies.

Finally, we show that our RDMA^{TSO} declarative semantics is equivalent to the simplified operational semantics in §3, with the full proof given in §D. Note that as a corollary (of theorem 3.1), our declarative semantics is also equivalent to the concrete operational semantics in §3.

THEOREM 4.5. *The RDMA^{TSO} declarative semantics is equivalent to the simplified operational one.*

Recall that RDMA^{TSO} assumes x86-TSO CPUs (subject to TSO consistency). As such, RDMA^{TSO} can be seen as an extension of TSO [Alglave et al. 2014]. That is, as we show in the following theorem, for programs without remote operations, RDMA^{TSO} and TSO coincide (see §C.3 for the full proof).

THEOREM 4.6. *For programs without remote operations, TSO- and RDMA^{TSO}-consistency coincide.*

Declarative Semantics without PCIe. To relax the PCIe-specific constraint (stipulating that NIC local/remote reads propagates all pending NIC writes in $\mathbf{wb_L}/\mathbf{wb_R}$ to memory), we adapt our declarative model as follows: (1) we no longer need the nfo relation in Def. 4.2 and Def. 4.3 (as nfo is used to capture the 'NIC flush order' under PCIe); (2) we replace sqp in cells G8 and G9 of the oppo table (Fig. 9) with ✗ (as a later nrR/nlW no longer flushes $\mathbf{wb_R}$ and thus may not observe an earlier nrW); and (3) we redefine $\mathsf{rf_b}$ and $\mathsf{rb_b}$ as follows to include events on the same queue pair:

$$\mathsf{rf_b} \triangleq ([\mathtt{lW}];(\mathsf{rf} \cap \mathsf{sthd});[\mathtt{lR}]) \cup (\mathsf{rf} \cap \mathsf{sqp}) \qquad\qquad \mathsf{rb_b} \triangleq ([\mathtt{lR}];(\mathsf{rb} \cap \mathsf{sthd});[\mathtt{lW}]) \cup (\mathsf{rb} \cap \mathsf{sqp})$$

Recall that $w \xrightarrow{\mathsf{rf_b}} r$ (resp. $r \xrightarrow{\mathsf{rb_b}} w$) denotes r reads from (resp. before) w before w is observable. In the strong model with PCIe guarantee, NIC local/remote reads flush the $\mathbf{wb_L}/\mathbf{wb_R}$, and thus r reads from (resp. before) w only once w is observable. By contrast, in the relaxed model without PCIe guarantee, NIC local/remote reads no longer flush the $\mathbf{wb_L}/\mathbf{wb_R}$, and thus r can read from (resp. before) w while it is still in $\mathbf{wb_L}/\mathbf{wb_R}$ (i.e. not yet observable). We thus expand the $\mathsf{rf_b}$ (resp. $\mathsf{rb_b}$) definition to account for the possibility of reading from (resp. before) a write before it is observable.

5 VALIDATING THE RDMA$^{\text{TSO}}$ MODEL

To complement our formal semantics, we conducted an extensive validation of litmus tests on two distinct setups: InfiniBand (IB) and RDMA over Converged Ethernet (RoCE). In the IB setup, we used Dell PowerEdge R740 x86-64 machines, Intel(R) Xeon(R) Gold 6132 CPUs (2.60GHz, 14 cores), Ubuntu 22.04.2 LTS with Mellanox Technologies MT28908 Family [ConnectX-6] controller. In the RoCE setup, we used machines with Intel(R) Xeon(R) E-2286G CPUs (4.00GHz, 14 cores), ran Linux kernel version 4.18.0, Red Hat Enterprise Linux version 8 (477.27.1), with a Mellanox Technologies MT27800 Family [ConnectX-5] controller. Each of our tests is written in C – see §A and the code in the supplementary material [Ambal et al. 2024].

Litmus Tests and Outcomes. We focused our validation efforts on a set of 37 litmus tests representative of a wide range of allowed and disallowed weak behaviours, including several variants of well-known concurrent tests in the literature such as 'store buffering' (SB), 'message passing' (MP), 'load buffering' (LB), 'independent reads of independent writes' (IRIW), parallel writes (2+2W) and so forth. Our results corroborate our RDMA$^{\text{TSO}}$ model and confirm that (1) RDMA$^{\text{TSO}}$ is *not too strong* (i.e. it does not prohibit behaviours exhibited by the hardware); and (2) RDMA$^{\text{TSO}}$ is *not too weak* (i.e. it does not admit too many weak behaviours not exhibited by hardware). Indeed, despite extensive testing we detected *no behaviours prohibited by RDMA^{TSO}*; and we observed *almost all behaviours allowed* by RDMA$^{\text{TSO}}$, with a few exceptions detailed below.

Bringing about Weak Behaviours. As with litmus testing for local (CPU-only) concurrency for established models such as TSO, observing a weak behaviour relies on several factors, including the order in which threads are scheduled and interleaved, as well as the timing of how writes are propagated and made visible to different threads (e.g. when writes are removed from the store buffer and propagated through the cache hierarchy and the memory). In the context of (RDMA) programs with remote operations, there are additional factors at play such as the NIC workload. As such, inducing weak behaviours necessitates creating suitable conditions, which may involve stressing the NIC and/or the CPU in the background. Indeed, as shown by the work of Dan et al. [2016], this is far from straightforward as they failed to observe *any* weak behaviours at all. To remedy this, we used several techniques to elicit the weak behaviours permitted by the specification. Specifically, since no testing tool currently exists for automatically overwhelming or decelerating an RDMA system, in close consultation with NVIDIA experts we devised several techniques for creating bottlenecks and fostering the emergence of weak behaviours, as outlined below:

- **NIC delays**: We introduced delays on the Network Interface Card (NIC) by initiating numerous flood RDMA read or write operations between the client and the server on the designated queue pair (QP) at a specific juncture within the litmus test. As these operations still obey the same ordering rules, they cause some of the other operations in the litmus test to be delayed until they are successfully completed.
- **CPU delays**: We strategically injected CPU delays of random length into the litmus tests at certain points. We did this by selecting a random duration from three possibilities: 0 nanoseconds, half of the *round-trip time* (RTT) and the full RTT, delaying the CPU execution by the chosen duration. RTT represents the round-trip time between the client and the server, capturing the duration it takes for a packet to travel from the sender to the receiver and back. We implemented the delay functionality using a busy-wait loop, where the CPU remains occupied, repeatedly polling the clock without executing additional tasks. Once the delay duration has elapsed, the loop concludes, allowing the CPU to resume normal execution. While such delays can bring about weak behaviours in certain tests, such as that in Fig. 2d or **LB2** in the technical appendix (Fig. 13 in §A.5), it may impede observing others such as that in Fig. 2b, which requires no delay between $z^2 := x$ and $x := 1$. Therefore, by incorporating random delays, we aimed to encompass the full spectrum of potential scenarios, ensuring comprehensive coverage of behavioural variations across different tests. Interestingly, we observed that even minimal delays introduced by printing variables for debugging purposes can disrupt the manifestation of behaviours in certain tests (e.g. that in Fig. 2b), hence we limited our use of print statements.
- **High RDMA traffic loads**: While introducing NIC delays as described above can sometimes help induce certain interleavings, e.g. by delaying one node while letting another to continue, in some cases delays do not help as both relevant operations are delayed by the same NIC delay operation. In such cases, using *background traffic* on unrelated queue pairs allows additional interleavings. Specifically, we generated concurrent high RDMA traffic loads on the background using numerous queue pairs and utilising the zero-copy mode (transferring data directly between the memory of the sender and receiver without intermediate copying). This increased RDMA activity introduces competition for shared system resources such as CPU cycles, memory bandwidth, network capacity and NIC processing capacity. This led to contention amongst threads in a litmus test, potentially altering their scheduling behaviour. This strategic approach proved pivotal in uncovering weak behaviours in several tests such as **MP3bis** in the technical appendix (see Fig. 11 in §A.3), which might otherwise have gone unnoticed. We observed that the number of queue pairs utilised to generate the traffic load plays a crucial role in certain tests. For instance, in the case of **GFP2** in Fig. 17 (§A.11 in the technical appendix), we used 128 queue pairs to induce the weak behaviour.
- **Synchronisation**: As is common practice in validation via litmus testing, we employed loops to bring about the desired weak behaviours. For instance, in order to ensure that a local (resp. remote) load operation reads the desired value required by a given weak behaviour, we place the operation within a loop (in effect replacing the single operation with multiple such operations), iterating until the desired value is read. In other cases, we placed the entire litmus test within a loop. Using loops this way can sometimes replace the CPU delay mechanism discussed above, allowing it to pinpoint the delay needed to reproduce the behaviour.

Validation Results. Using the techniques above, we successfully confirmed almost all weak behaviours allowed by RDMA$^{\text{TSO}}$, including all ✓ examples in Figs. 1 to 3, as well as several variants of MP, SB, LB, IRIW, 2+2W, and others, observed at varying frequencies (see §A for more details). The frequency of observing these behaviours is influenced not only by the nuances of the litmus test itself, but also by the infrastructure settings such as the network setup, switch configuration and NIC capabilities. Moreover, dynamic and unpredictable factors such as network congestion, packet drops, latency fluctuations, bandwidth utilisation and routing changes also significantly impact the observations. An interesting phenomenon we encountered was the variability in behaviour manifestation even within a specific test scenario. For example, in the case of **GFP2** (Fig. 17, §A.11 in the technical appendix), the weak behaviour was observed at one of our test premises but not the other (we had two test premises at two distinct geographical locations). This observation highlights the intricacy and subtlety involved in inducing weak behaviours.

In 4 of the litmus tests, we did not observe the allowed weak behaviour because the hardware implementation did not utilise the weakness permitted by the specification. As discussed in §2, manifesting the weak behaviour in such cases relies on polling a put operation before the remote NIC write completes, e.g. as in Fig. 3d. This weak behaviour is allowed by RDMA^{TSO} because the specification allows the remote NIC to return an acknowledgement *before* performing the associated write. Nevertheless, according to NVIDIA experts, this weakness is not utilised by the hardware implementations. To the best of our knowledge, these weak behaviours cannot be observed on current hardware implementations, but they might emerge in future ones.

Limitations. The main limitation of our validation is that weak behaviours are only exposed by hand-crafted techniques (discussed above) that stress the system in certain ways. This is currently a difficult trial-and-error process that requires a high degree of knowledge of current hardware implementations (which we acquired through close consultation with NVIDIA experts). As such, executing RDMA tests is currently not amenable to *mass automation* as in the frameworks of [Alglave et al. 2021, 2014; Raad et al. 2022]. To adapt these frameworks for RDMA, one would need to develop systematic heuristics for automatically applying the techniques discussed above. We leave this challenge to future work.

6 RELATED AND FUTURE WORK

RDMA Semantics. To our knowledge, the only existing work on the formal semantics of RDMA programs is that of coreRMA [Dan et al. 2016], which has several key limitations, as follows. (1) Although Dan et al. [2016] attempted to validate coreRMA, they observed *none* of the weak behaviours allowed by coreRMA in existing hardware implementations. This is in contrast to our work, where we have observed almost all weak behaviours allowed by RDMA^{TSO} in existing implementations, and in the rare cases where we could not observe a behaviour, we have confirmed that this is because existing implementations explicitly did not utilise the weakness allowed by the specification. (2) coreRMA assumes that CPU concurrency is governed by the strong and unrealistic SC model [Lamport 1979]. (3) coreRMA only presents a declarative (and not operational) characterisation. (4) Most importantly, coreRMA departs from the RDMA specification [IBTA 2022] in three important ways, as follows.

First, they do not model the poll operation, $\texttt{poll}(n)$, described in the specification. Instead, they model a flush operation, $\texttt{flush}(n)$, whose behaviour does not match any operation defined in the specification. Specifically, $\texttt{flush}(n)$ waits for all previous RDMA operations on the n-queue pair to complete and further blocks later CPU operations and RDMA operations on the *same* queue pair. In other words, $\texttt{flush}(n)$ is tantamount to $\texttt{rfence}(n)$ followed by $\texttt{mfence}$. However, $\texttt{flush}(n)$ does *not* block later RDMA operations on *different* queue pairs, and hence under coreRMA there is no clean way to enforce an order on two RDMA operations on different queue pairs.

For instance, to ensure that $x := y^2; z^3 := x$ is executed in order (so that z gets the updated value of x), it is not enough to add a flush and write $x := y^2; \underline{\texttt{flush(2)}}; z^3 := x$. Instead, one must also add a superfluous *CPU operation* and write e.g. $x := y^2; \underline{\texttt{flush}(2); a := w}; z^3 := x$. This is because $\texttt{flush}(2)$ blocks the CPU operation $a := w$, which in turn blocks $z^3 := x$.

Second, RDMA$^{\text{TSO}}$ preserves the order between two NIC local reads and two NIC local writes (cells F6 and I9 in Fig. 9), while coreRMA does not, violating the RDMA specification [IBTA 2022]. For instance, given $y^2 := x; z^2 := w$, RDMA$^{\text{TSO}}$ guarantees x is read before w, while there is no such guarantee under coreRMA; i.e. in such scenarios coreRMA is *weaker* than the specification.

Finally, as per the RDMA specification, under RDMA$^{\text{TSO}}$ remote reads can be reordered with respect to other remote operations on the same queue pair (cells H7 and H8), while under coreRMA they cannot; i.e. in such scenarios coreRMA is *stronger* than the specification. Consequently, as coreRMA is weaker than the specification in some scenarios and stronger in others, it is *neither a strict abstraction, nor a strict refinement* of the specification.

Weak Memory Models. Existing literature includes several examples of weak consistency models, both at hardware and software levels. On the hardware side, several works have formalised the semantics of the x86 architecture [Abdulla et al. 2015; Alglave et al. 2014; Raad et al. 2022; Sewell et al. 2010]. However, none of these works covered the consistency semantics of RDMA programs in the context of x86 machines. Similarly, several works have formalised the semantics of the ARMv8 and POWER architectures, both operationally and declaratively [Alglave et al. 2021, 2014; Chakraborty and Vafeiadis 2019; Flur et al. 2016; Mador-Haim et al. 2012; Pulte et al. 2018; Sarkar et al. 2011]. On the software side, there has been a number of formal models for C11 consistency [Batty et al. 2011; Kang et al. 2017; Lahav et al. 2016, 2017; Lee et al. 2020; Nienhuis et al. 2016; Pichon-Pharabod and Sewell 2016] with verified compilation schemes [Moiseenko et al. 2020; Podkopaev et al. 2017, 2019], Java [Bender and Palsberg 2019; Manson et al. 2005], transactional memory [Raad et al. 2018, 2019a; Xiong et al. 2020], the Linux kernel [Alglave et al. 2018] and the ext4 filesystem [Kokologiannakis et al. 2021]. Additionally, there has been several works on formalising the *persistency* semantics of programs in the context of non-volatile memory, describing the behaviour of programs in case of crashes [Cho et al. 2021; Khyzha and Lahav 2021; Raad and Vafeiadis 2018; Raad et al. 2020b, 2019b], as well as program logics for verifying such programs [Bila et al. 2022; Raad et al. 2020a].

Future Work. We plan to build over our work presented here in several ways. First, we aim to adapt existing methods for automatically generating litmus tests (e.g. [Alglave et al. 2021, 2014]) to RDMA programs by developing heuristics for automatically applying the inducement/stressing techniques discussed in §5 for bringing about weak behaviours. Second, we plan to formalise the semantics of RDMA programs in the context of the ARMv8 hardware architecture (i.e. when CPU concurrency is governed by ARMv8 rather than TSO). Third, we plan to verify existing RDMA implementations such as those of the Verbs [linux-rdma 2018] and libfabric [OpenFabrics 2016] APIs. Lastly, using our formal RDMA$^{\text{TSO}}$ semantics, we plan to develop *manual* reasoning techniques such as program logics (underpinned by our operational semantics), as well as *automated* verification techniques such as model checking (based on our declarative semantics) for RDMA.

ACKNOWLEDGMENTS

We would like to thank Alexey Gotsman, Adam Morrison, Noam Rinetzky, and especially Yuri Meshman, for starting this research and whose preliminary memory model underpins the results presented here. Additional thanks to Yamin Friedman, Daniel Marcovitch and Liran Liss for sharing their insights into the IBTA specifications and NVIDIA's RDMA implementations. We also thank the anonymous reviewers for their valuable feedback and Viktor Vafeiadis for many fruitful discussions. Guillaume Ambal is supported by the EPSRC grant EP/X037029/1. Brijesh Dongol is supported by VeTSS and EPSRC projects EP/Y036425/1, EP/X037142/1, EP/X015149/1, EP/V038915/1, and EP/R025134/2. Ori Lahav is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 851811) and the Israel Science Foundation (grant number 814/22). Azalea Raad is supported by a UKRI fellowship MR/V024299/1, by the EPSRC grant EP/X037029/1, and by VeTSS.

REFERENCES

Parosh Aziz Abdulla, Mohamed Faouzi Atig, Ahmed Bouajjani, K. Narayan Kumar, and Prakash Saivasan. 2021. Deciding Reachability under Persistent X86-TSO. *Proc. ACM Program. Lang.* 5, POPL, Article 56 (Jan. 2021), 32 pages. https://doi.org/10.1145/3434337

Parosh Aziz Abdulla, Mohamed Faouzi Atig, and Tuan-Phong Ngo. 2015. The Best of Both Worlds: Trading Efficiency and Optimality in Fence Insertion for TSO. In *Proceedings of the 24th European Symposium on Programming on Programming Languages and Systems - Volume 9032*. Springer-Verlag New York, Inc., New York, NY, USA, 308–332. https://doi.org/10.1007/978-3-662-46669-8_13

Marcos K. Aguilera, Naama Ben-David, Rachid Guerraoui, Virendra J. Marathe, and Igor Zablotchi. 2019. The Impact of RDMA on Agreement. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, Peter Robinson and Faith Ellen (Eds.). ACM, 409–418. https://doi.org/10.1145/3293611.3331601

Jade Alglave, Will Deacon, Richard Grisenthwaite, Antoine Hacquard, and Luc Maranget. 2021. Armed Cats: Formal Concurrency Modelling at Arm. *ACM Trans. Program. Lang. Syst.* 43, 2 (2021), 8:1–8:54. https://doi.org/10.1145/3458926

Jade Alglave, Luc Maranget, Paul E. McKenney, Andrea Parri, and Alan Stern. 2018. Frightening Small Children and Disconcerting Grown-Ups: Concurrency in the Linux Kernel. *SIGPLAN Not.* 53, 2 (March 2018), 405–418. https://doi.org/10.1145/3296957.3177156

Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory. *ACM Trans. Program. Lang. Syst.* 36, 2 (2014), 7:1–7:74. https://doi.org/10.1145/2627752

Guillaume Ambal, Brijesh Dongol, Haggai Eran, Vasileios Klimis, Ori Lahav, and Azalea Raad. 2024. Project page for Semantics of Remote Direct Memory Access. https://www.soundandcomplete.org/papers/OOPSLA2024/RDMA

Don Anderson. 1999. *FireWire system architecture (2nd ed.): IEEE 1394a*. Addison-Wesley Longman Publishing Co., Inc., USA.

Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. 2011. Mathematizing C++ Concurrency. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Austin, Texas, USA) *(POPL '11)*. ACM, New York, NY, USA, 55–66. https://doi.org/10.1145/1926385.1926394

John Bender and Jens Palsberg. 2019. A Formalization of Java's Concurrent Access Modes. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 142 (Oct. 2019), 28 pages. https://doi.org/10.1145/3360568

Eleni Vafeiadi Bila, Brijesh Dongol, Ori Lahav, Azalea Raad, and John Wickerson. 2022. View-Based Owicki–Gries Reasoning for Persistent x86-TSO. In *Programming Languages and Systems*, Ilya Sergey (Ed.). Springer International Publishing, Cham, 234–261.

M. S. Birrittella, M. Debbage, R. Huggahalli, J. Kunz, T. Lovett, T. Rimmer, K. D. Underwood, and R. C. Zak. 2015. Intel Omni-Path Architecture: Enabling scalable, high performance fabrics. In *2015 IEEE 23rd Annual Symposium on High-Performance Interconnects (HOTI) (HOTI 2015)*. 1–9. https://doi.org/10.1109/HOTI.2015.22

Ahmed Bouajjani, Egor Derevenetc, and Roland Meyer. 2013. Checking and Enforcing Robustness against TSO. In *ESOP 2013 (LNCS, Vol. 7792)*. Springer, 533–553. https://doi.org/10.1007/978-3-642-37036-6_29

Soham Chakraborty and Viktor Vafeiadis. 2019. Grounding Thin-Air Reads with Event Structures. *Proc. ACM Program. Lang.* 3, POPL, Article 70 (Jan. 2019), 28 pages. https://doi.org/10.1145/3290383

Kyeongmin Cho, Sung-Hwan Lee, Azalea Raad, and Jeehoon Kang. 2021. Revamping Hardware Persistency Models: View-Based and Axiomatic Persistency Models for Intel-X86 and Armv8. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (Virtual, Canada) *(PLDI 2021)*. Association for Computing Machinery, New York, NY, USA, 16–31. https://doi.org/10.1145/3453483.3454027

Andrei Marian Dan, Patrick Lam, Torsten Hoefler, and Martin Vechev. 2016. Modeling and Analysis of Remote Memory Access Programming. *SIGPLAN Not.* 51, 10 (oct 2016), 129–144. https://doi.org/10.1145/3022671.2984033

D. Dunning, G. Regnier, G. McAlpine, D. Cameron, B. Shubert, F. Berry, A.M. Merritt, E. Gronke, and C. Dodd. 1998. The Virtual Interface Architecture. *IEEE Micro* 18, 2 (1998), 66–76. https://doi.org/10.1109/40.671404

Shaked Flur, Kathryn E. Gray, Christopher Pulte, Susmit Sarkar, Ali Sezgin, Luc Maranget, Will Deacon, and Peter Sewell. 2016. Modelling the ARMv8 Architecture, Operationally: Concurrency and ISA. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (St. Petersburg, FL, USA) *(POPL '16)*. Association for Computing Machinery, New York, NY, USA, 608–621. https://doi.org/10.1145/2837614.2837615

Robert Gerstenberger, Maciej Besta, and Torsten Hoefler. 2018. Enabling Highly Scalable Remote Memory Access Programming with MPI-3 One Sided. *Commun. ACM* 61, 10 (sep 2018), 106–113. https://doi.org/10.1145/3264413

IBTA. 2022. InfiniBand Architecture Specification Volume 1 Release 1.6. https://www.infinibandta.org/ibta-specification/.

InfiniBand Trade Association (IBTA). 2018. The RoCE Initiative. https://www.infinibandta.org/roce-initiative/ (Accessed: July 2023).

Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A Promising Semantics for Relaxed-Memory Concurrency. *SIGPLAN Not.* 52, 1 (Jan. 2017), 175–189. https://doi.org/10.1145/3093333.3009850

Artem Khyzha and Ori Lahav. 2021. Taming X86-TSO Persistency. *Proc. ACM Program. Lang.* 5, POPL, Article 47 (Jan. 2021), 29 pages. https://doi.org/10.1145/3434328

Michalis Kokologiannakis, Ilya Kaysin, Azalea Raad, and Viktor Vafeiadis. 2021. PerSeVerE: Persistency Semantics for Verification under Ext4. *Proc. ACM Program. Lang.* 5, POPL, Article 43 (jan 2021), 29 pages. https://doi.org/10.1145/3434324

Ori Lahav and Udi Boker. 2020. Decidable verification under a causally consistent shared memory. In *PLDI 2020*, Alastair F. Donaldson and Emina Torlak (Eds.). ACM, 211–226. https://doi.org/10.1145/3385412.3385966

Ori Lahav, Nick Giannarakis, and Viktor Vafeiadis. 2016. Taming Release-Acquire Consistency. *SIGPLAN Not.* 51, 1 (Jan. 2016), 649–662. https://doi.org/10.1145/2914770.2837643

Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. Repairing Sequential Consistency in C/C++11. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Barcelona, Spain) *(PLDI 2017)*. Association for Computing Machinery, New York, NY, USA, 618–632. https://doi.org/10.1145/3062341.3062352

Leslie Lamport. 1979. How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs. *IEEE Trans. Computers* 28, 9 (Sept. 1979), 690–691. https://doi.org/10.1109/TC.1979.1675439

Sung-Hwan Lee, Minki Cho, Anton Podkopaev, Soham Chakraborty, Chung-Kil Hur, Ori Lahav, and Viktor Vafeiadis. 2020. Promising 2.0: Global Optimizations in Relaxed Memory Concurrency. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation* (London, UK) *(PLDI 2020)*. Association for Computing Machinery, New York, NY, USA, 362–376. https://doi.org/10.1145/3385412.3386010

linux-rdma. 2018. RDMA core. https://github.com/linux-rdma/rdma-core/ (Accessed: Jul. 2023).

Sela Mador-Haim, Luc Maranget, Susmit Sarkar, Kayvan Memarian, Jade Alglave, Scott Owens, Rajeev Alur, Milo M. K. Martin, Peter Sewell, and Derek Williams. 2012. An Axiomatic Memory Model for POWER Multiprocessors. In *Computer Aided Verification - 24th International Conference, CAV 2012, Berkeley, CA, USA, July 7-13, 2012 Proceedings (Lecture Notes in Computer Science, Vol. 7358)*, P. Madhusudan and Sanjit A. Seshia (Eds.). Springer, 495–512. https://doi.org/10.1007/978-3-642-31424-7_36

Jeremy Manson, William Pugh, and Sarita V. Adve. 2005. The Java Memory Model. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Long Beach, California, USA) *(POPL '05)*. Association for Computing Machinery, New York, NY, USA, 378–391. https://doi.org/10.1145/1040305.1040336

Evgenii Moiseenko, Anton Podkopaev, Ori Lahav, Orestis Melkonian, and Viktor Vafeiadis. 2020. Reconciling Event Structures with Modern Multiprocessors (Artifact). *Dagstuhl Artifacts Series* 6, 2 (2020), 4:1–4:3. https://doi.org/10.4230/DARTS.6.2.4

Kyndylan Nienhuis, Kayvan Memarian, and Peter Sewell. 2016. An Operational Semantics for C/C++11 Concurrency. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications* (Amsterdam, Netherlands) *(OOPSLA 2016)*. Association for Computing Machinery, New York, NY, USA, 111–128. https://doi.org/10.1145/2983990.2983997

Stanko Novakovic, Alexandros Daglis, Edouard Bugnion, Babak Falsafi, and Boris Grot. 2014. Scale-out NUMA. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems* (Salt Lake City, Utah, USA) *(ASPLOS '14)*. Association for Computing Machinery, New York, NY, USA, 3–18. https://doi.org/10.1145/2541940.2541965

NVIDIA Corporation. 2021. NVIDIA BlueField-2 DPU. https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/documents/datasheet-nvidia-bluefield-2-dpu.pdf (Accessed: Jul. 2023).

OpenFabrics. 2016. RDMA core. https://ofiwg.github.io/libfabric/ (Accessed: Jul. 2023).

PCI-SIG. 2022. PCI Express Base Specification Revision 6.0 Version 1.0. https://pcisig.com/pci-express-6.0-specification.

Jean Pichon-Pharabod and Peter Sewell. 2016. A Concurrency Semantics for Relaxed Atomics That Permits Optimisation and Avoids Thin-Air Executions. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (St. Petersburg, FL, USA) *(POPL '16)*. Association for Computing Machinery, New York, NY, USA, 622–633. https://doi.org/10.1145/2837614.2837616

Anton Podkopaev, Ori Lahav, and Viktor Vafeiadis. 2017. Promising Compilation to ARMv8 POP. In *31st European Conference on Object-Oriented Programming (ECOOP 2017) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 74)*, Peter Müller (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 22:1–22:28. https://doi.org/10.4230/LIPIcs.ECOOP.2017.22

Anton Podkopaev, Ori Lahav, and Viktor Vafeiadis. 2019. Bridging the Gap Between Programming Languages and Hardware Weak Memory Models. *Proc. ACM Program. Lang.* 3, POPL, Article 69 (Jan. 2019), 31 pages. https://doi.org/10.1145/3290382

Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. 2018. Simplifying ARM Concurrency: Multicopy-atomic Axiomatic and Operational Models for ARMv8. *Proc. ACM Program. Lang.* 2, POPL, Article 19 (Dec. 2018), 29 pages. https://doi.org/10.1145/3158107

Christopher Pulte, Jean Pichon-Pharabod, Jeehoon Kang, Sung-Hwan Lee, and Chung-Kil Hur. 2019. Promising-ARM/RISC-V: A Simpler and Faster Operational Concurrency Model. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Phoenix, AZ, USA) *(PLDI 2019)*. Association for Computing Machinery, New York, NY, USA, 1–15. https://doi.org/10.1145/3314221.3314624

Azalea Raad, Ori Lahav, and Viktor Vafeiadis. 2018. On Parallel Snapshot Isolation and Release/Acquire Consistency. In *Programming Languages and Systems*, Amal Ahmed (Ed.). Springer International Publishing, Cham, 940–967.

Azalea Raad, Ori Lahav, and Viktor Vafeiadis. 2019a. On the Semantics of Snapshot Isolation. In *Verification, Model Checking, and Abstract Interpretation*, Constantin Enea and Ruzica Piskac (Eds.). Springer International Publishing, Cham, 1–23.

Azalea Raad, Ori Lahav, and Viktor Vafeiadis. 2020a. Persistent Owicki-Gries Reasoning: A Program Logic for Reasoning about Persistent Programs on Intel-X86. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 151 (nov 2020), 28 pages. https://doi.org/10.1145/3428219

Azalea Raad, Luc Maranget, and Viktor Vafeiadis. 2022. Extending Intel-X86 Consistency and Persistency: Formalising the Semantics of Intel-X86 Memory Types and Non-Temporal Stores. *Proc. ACM Program. Lang.* 6, POPL, Article 22 (jan 2022), 31 pages. https://doi.org/10.1145/3498683

Azalea Raad and Viktor Vafeiadis. 2018. Persistence Semantics for Weak Memory: Integrating Epoch Persistency with the TSO Memory Model. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 137 (Oct. 2018), 27 pages. https://doi.org/10.1145/3276507

Azalea Raad, John Wickerson, Gil Neiger, and Viktor Vafeiadis. 2020b. Persistency Semantics of the Intel-X86 Architecture. *Proc. ACM Program. Lang.* 4, POPL, Article 11 (Dec. 2020), 31 pages. https://doi.org/10.1145/3371079

Azalea Raad, John Wickerson, and Viktor Vafeiadis. 2019b. Weak Persistency Semantics from the Ground Up: Formalising the Persistency Semantics of ARMv8 and Transactional Models. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 135 (Oct. 2019), 27 pages. https://doi.org/10.1145/3360561

Renato J. Recio, Paul R. Culley, Dave Garcia, Bernard Metzler, and Jeff Hilland. 2007. A Remote Direct Memory Access Protocol Specification. RFC 5040. https://doi.org/10.17487/RFC5040

Susmit Sarkar, Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams. 2011. Understanding POWER Multiprocessors. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation* (San Jose, California, USA) *(PLDI '11)*. Association for Computing Machinery, New York, NY, USA, 175–186. https://doi.org/10.1145/1993498.1993520

Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. 2010. X86-TSO: A Rigorous and Usable Programmer's Model for x86 Multiprocessors. *Commun. ACM* 53, 7 (July 2010), 89–97. https://doi.org/10.1145/1785414.1785443

Alexander Shpiner, Eitan Zahavi, Omar Dahley, Aviv Barnea, Rotem Damsker, Gennady Yekelis, Michael Zus, Eitan Kuta, and Dean Baram. 2017. RoCE Rocks without PFC: Detailed Evaluation. In *Proceedings of the Workshop on Kernel-Bypass Networks* (Los Angeles, CA, USA) *(KBNets '17)*. Association for Computing Machinery, New York, NY, USA, 25–30. https://doi.org/10.1145/3098583.3098588

SPARC. 1992. *The SPARC Architecture Manual: Version 8.* Prentice-Hall, Inc., Upper Saddle River, NJ, USA.

S. Van Doren. 2019. Abstract - HOTI 2019: Compute Express Link. In *2019 IEEE Symposium on High-Performance Interconnects (HOTI) (HOTI 2019)*. 18–18. https://doi.org/10.1109/HOTI.2019.00017

Xingda Wei, Jiaxin Shi, Yanzhe Chen, Rong Chen, and Haibo Chen. 2015. Fast In-Memory Transaction Processing Using RDMA and HTM. In *Proceedings of the 25th Symposium on Operating Systems Principles* (Monterey, California) *(SOSP '15)*. Association for Computing Machinery, New York, NY, USA, 87–104. https://doi.org/10.1145/2815400.2815419

Shale Xiong, Andrea Cerone, Azalea Raad, and Philippa Gardner. 2020. Data Consistency in Transactional Storage Systems: A Centralised Semantics. In *34th European Conference on Object-Oriented Programming (ECOOP 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 166)*, Robert Hirschfeld and Tobias Pape (Eds.). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 21:1–21:31. https://doi.org/10.4230/LIPIcs.ECOOP.2020.21

Yibo Zhu, Haggai Eran, Daniel Firestone, Chuanxiong Guo, Marina Lipshteyn, Yehonatan Liron, Jitendra Padhye, Shachar Raindel, Mohamad Haj Yahia, and Ming Zhang. 2015. Congestion Control for Large-Scale RDMA Deployments. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication* (London, United Kingdom) *(SIGCOMM '15)*. Association for Computing Machinery, New York, NY, USA, 523–536. https://doi.org/10.1145/2785956.2787484

A LITMUS TESTS

In this section, we present a wide range of RDMA tests exhibiting different weak behaviours. We always use at least two nodes, and do not discuss tests on CPU concurrency (such as the one of Fig. 1) as they simply follow the known behaviours of the usual TSO concurrency.

By convention, a, b, c, and d can be seen as local locations and are not used by other threads. There are always initialised to 0, even if not mentioned. We use names x, y, z, v, and w for locations accessed by multiple threads.

In each example, we number threads T1, T2, T3, etc., and nodes N1, N2, etc., from left to right.

For each litmus test, we use the following check-marks to annotate the given behaviours:

- ✗ The behaviour is *not* allowed by the specification, and effectively we did *not* observe it in our testing.
- ✓ The behaviour *is* allowed by the specification, and we *did* observe it in our testing.
- ✓* The behaviour *is* allowed by the specification, but we did *not* observe it in out testing. In most cases, this is because modern hardware and setups do *not* allow it.

A.1 Limitations of Experimental Behaviours

While the RDMA semantics specification allows for a wide range of weak behaviours, current hardware and setups do not take full advantage of this weakness.

The main offender is when polling a put operation (see e.g. test **MP3**). The semantics allows for the write to be delayed after the return of the completion notification. In practice, $\mathbf{wb_R}$ is part of the PCIe domain and the write cannot be reordered with respect to other instructions, such as completion notifications of queue pairs of the remote node. As such, we do not expect to observe any behaviour using this feature of the specification.

A.2 Single-Threaded

ST1

$x=0$	$z=0$
$x := 1$ $z^2 := x$	

$z = 0$ ✗

ST2

$x=0$	$z=0$
$z^2 := x$ $x := 1$	

$z = 1$ ✓

ST3

$x=0$	$z=0$
$z^2 := x$ poll(2) $x := 1$	

$z = 1$ ✗

ST4

$x=0$	$z=0$
$z^2 := x$ $z^2 := x$ poll(2) $x := 1$	

$z = 1$ ✓

ST5

$x=0$	$z=1$	$y=2$
$x := z^2$ $x := y^3$		

$x = 1$ ✓

ST6

$x=0$	$z=1$	$y=2$
$y^3 := x$ $x := z^2$		

$y = 0$ ✓

ST7

$x=1$	$y=z=0$
$z^2 := x$ $x := y^2$	

$z = 0$ ✗

ST8

$x = 1$	$y=z=0$
$x := y^2$ $z^2 := x$	

$z = 1$ ✓

ST9

$x = 1$	$y=z=0$
$x := y^2$ rfence (2) $z^2 := x$	

$z = 1$ ✗

Fig. 10. Single-Threaded Litmus Tests.

Figure 10 contains a copy of the single-threaded examples of Fig. 2. All allowed behaviours are observable in practice.

A.3 Message Passing

The examples are presented in Fig. 11.

MP1	
	$x, y = 0$
$x^2 := 1$ $y^2 := 1$	$a := y$ $b := x$

$(a, b) = (1, 0)$ ✗

MP2	
$x = 0$	$y = 0$
$x := 1$ $y^2 := 1$	$a := y$ $b := x^1$

$(a, b) = (1, 0)$ ✗

MP3	
$y = 0$	$x = 0$
$x^2 := 1$ `poll(2)` $y := 1$	$a := y^1$ `poll(1)` $b := x$

$(a, b) = (1, 0)$ ✓*

MP3bis	
$y = 0$	$x = 0$
$x^2 := 1$ $y := 1$	$a := y^1$ $b := x$

$(a, b) = (1, 0)$ ✓

MP4	
$x, y = 0$	
$x := 1$ $y := 1$	$a := y^1$ $b := x^1$

$(a, b) = (1, 0)$ ✓*

MP4bis	
$x, y = 0$	
$x := 1$ $y := 1$	$a := y^1$ `rfence (1)` $b := x^1$

$(a, b) = (1, 0)$ ✗

Fig. 11. Message Passing Litmus Tests.

If x is on T1 and y on T2, the weak behaviour is not possible: both $(x := 1)$ and $(a := y)$ are local, so one has to finish first before the remote operations can start. If both are in T2, local read order and remote write order (on the same queue pair) are preserved so this is not possible.

If x is on T2 and y on T1 we can observe this behaviour (**MP3bis**) as remote operations can be delayed after local ones. Following the semantics of the paper, even `poll` instructions cannot prevent this (see test **MP3**), as they do not force a flush of the $\mathbf{wb_R}$ buffer. We can block $(x^2 := 1)$ in $\mathbf{wb_R}$ before executing $(y := 1)$, $(a := y^1)$, and $(b := x)$ in order. In practice, we do not expect **MP3** to be observable.

This behaviour is also allowed for x and y both on T1 (test **MP4**), because we can reorder remote reads. We first send both $(a := y^1)$ and $(b := x^1)$ to $\mathbf{out_R}$. Next, we read the value of x to transform the second instruction to $(b := 0)$. Then we fully execute $(x := 1)$ and $(y := 1)$. We can finally read y and execute $(a := 1)$ and $(b := 0)$. Note that, while we believe this behaviour should be possible in practice, we did not yet observe it in our testing. This would require manipulating the packets of the network and create a race between the answer packets and the resend of the work requests.

This behaviour can be prevented with an `rfence` or a `poll` instruction between the two remote operations (e.g., **MP4bis**).

A.4 Store Buffering

The examples are in Fig. 12.

If x is on T1 and y is on T2, this is not possible: both writing on x and y are local and one has to finish first before the remote reads can start.

If x and y are on the same node, the remote read on x flushes the remote write on y, so an `mfence` can prevent the behaviour. However, the weak behaviour is possible without `mfence`.

If x is on T2 and y is on T1, even `poll` instructions are not enough to prevent this, as they do not flush remote writes. Once again, we do not expect **SB3** to be observable in practice.

A.5 Load Buffering

The examples are in Fig. 13.

SB1

$x = 0$	$y = 0$
$x := 1$ $a := y^2$	$y := 1$ $b := x^1$

$(a, b) = (0, 0)$ ✗

SB2

$x, y = 0$	
$x := 1$ `mfence` $a := y$	$y^1 := 1$ $b := x^1$

$(a, b) = (0, 0)$ ✗
same for the symmetric

SB3

$y = 0$	$x = 0$
$x^2 := 1$ `poll(2)` $a := y$	$y^1 := 1$ `poll(1)` $b := x$

$(a, b) = (0, 0)$ ✓*

SB3bis

$y = 0$	$x = 0$
$x^2 := 1$ $a := y$	$y^1 := 1$ $b := x$

$(a, b) = (0, 0)$ ✓

Fig. 12. Store Buffering Litmus Tests.

LB1

$y = 0$	$x = 0$
$a := y$ $x^2 := 1$	$b := x$ $y^1 := 1$

$(a, b) = (1, 1)$ ✗

LB2

$x, y = 0$	
$a := y$ $x := 1$	$b := x^1$ $y^1 := 1$

$(a, b) = (1, 1)$ ✓
same for the symmetric

LB3

$x = 0$	$y = 0$
$a := y^2$ $x := 1$	$b := x^1$ $y := 1$

$(a, b) = (1, 1)$ ✓

LB3bis

$x = 0$	$y = 0$
$a := y^2$ `poll(2)` $x := 1$	$b := x^1$ `poll(1)` $y := 1$

$(a, b) = (1, 1)$ ✗

Fig. 13. Load Buffering Litmus Tests.

If x is on T2 and y is on T1, this is not possible: both reads on x and y are local and one has to finish first before the remote reads can start. The other 3 cases can present this behaviour.

If x and y are on the same process (wlog, let us assume T1), we can launch the get $(b := x^1)$ and make it wait in the $\mathbf{out_R}$ buffer. Then we can make $(y^1 := 1)$, $(a := y)$, and $(x := 1)$ fully execute in order, before finally reading $(x = 1)$ to write $(b = 1)$. This happens because remote write can take effect before a previous remote read. Note that adding an `rfence` (1) instruction on T2 would prevent this behaviour.

If x is on T1 and y on T2, we can proceed similarly. We can even observe this behaviour with one `poll` instruction. We make $(b := x^1)$ wait in $\mathbf{out_R}$, which does not prevent the local instruction $(y := 1)$ from executing, and we can finish as previously. Using two `poll` instructions would prevent the behaviour.

A.6 Independent Reads of Independent Writes

The threads T2 and T3 will always agree on the values of x and y. To obtain this behaviour, we thus need to reorder the read operations. This can be done as long as one of T2/T3 makes a remote read followed by a local read. The strategy is similar to the one for Load Buffering above, and it can resist one `poll` instruction. See **IRIW1** in Fig. 14 as an example.

We start with $(c := x^1)$ and let the operation stuck in $\mathbf{out_R}$. Then we fully execute $(d := y)$, $(y := 1)$, $(a := y^2)$, $(b := x)$, and $(x := 1)$. Finally, we can terminate the operation $(c := x^1)$ and get the desired behaviour. Similarly to previous tests, placing `poll` operations between remote and local operations would prevent this behaviour (see **IRIW2**).

A.7 2+2W

The examples are in Fig. 15.

IRIW1			
$x = 0$		$y = 0$	
$x := 1$	$a := y^2$ $b := x$	$c := x^1$ $d := y$	$y := 1$

$(a, b, c, d) = (1, 0, 1, 0)$ ✓

IRIW2			
$x = 0$		$y = 0$	
$x := 1$	$a := y^2$ `poll(2)` $b := x$	$c := x^1$ `poll(1)` $d := y$	$y := 1$

$(a, b, c, d) = (1, 0, 1, 0)$ ✗

Fig. 14. IRIW Litmus Tests.

2+2W1	
$x, y = 0$	
$x := 2$ $y := 1$	$y^1 := 2$ $x^1 := 1$

$(x, y) = (2, 2)$ ✗
same for the symmetric

2+2W2	
$x = 0$	$y = 0$
$x := 2$ $y^2 := 1$	$y := 2$ $x^1 := 1$

$(x, y) = (2, 2)$ ✗

2+2W3	
$y = 0$	$x = 0$
$x^2 := 2$ `poll(2)` $y := 1$	$y^1 := 2$ `poll(1)` $x := 1$

$(x, y) = (2, 2)$ ✓*

2+2W3bis	
$y = 0$	$x = 0$
$x^2 := 2$ $y := 1$	$y^1 := 2$ $x := 1$

$(x, y) = (2, 2)$ ✓

Fig. 15. 2+2 Writes Litmus Tests.

Note that the tests assume that everything that can be executed is executed. For instance, on the first test **2+2W1**, we could have $(x = 2 \wedge y = 2)$ for a long time if the remote write $x^1 := 1$ takes time to leave $\mathbf{wb_R}$ and be committed to memory.

If x and y are on the same thread, this cannot happen because both local write order and remote write order (on the same queue pair) are preserved. If x is on T1 and y is on T2, this cannot happen: one of the local writes ($\_ := 2$) has to finish before the remote operations can start.

If x is on T2 and y on T1, the behaviour can be observed, even with two `poll` instructions. Once again, polling does not force a message to arrive, so both remote writes can get stuck in the $\mathbf{wb_R}$ buffers long enough for both local writes to take effect. We expect **2+2W3bis** to be observable in practice, but not **2+2W3**.

A.8 Reordering different QP same node

In the semantics presented in this paper, each thread uses its own queue pair for each remote node. The semantics restricts reordering within a single queue pair, but does not limit reorderings on two queue pairs going from and to the same nodes.

QP1		
$y = 0$		$x, z = 0$
$x^2 := 1$ `poll(2)` $y := 1$	$a := y$ $z^2 := 1$	$b := z$ $c := x$

$(a, b, c) = (1, 1, 0)$ ✓*

In this example, if $(a = 1)$ at the end, then the operation $(x^2 := 1)$ was sent to a $\mathbf{wb_R}$ buffer on N2 before $(z^2 := 1)$ was sent to the beginning of the (T2/N2) queue pair. Then, $(b, c) = (1, 0)$ asserts that $(z^2 := 1)$ took effect before $(x^2 := 1)$, which is allowed by our semantics.

In practice, the (T1/N2) and (T2/N2) queue pairs would very likely use the same NIC and the operations of the two $\mathbf{wb_R}$ buffers cannot be reordered, which would prevent this behaviour allowed by the specification.

A.9 Read From Future

A remote write can overtake a remote read. But it then has to overtake all previous pending remote reads. If a get looks up a value from a future put, then other gets executed after have to consider puts up to this future.

RFF1	
$z = 0$	$x, y = 0$
$z := x^2$ $a := y^2$ $y^2 := 2$ $x^2 := 1$	$y := z^1$ `poll(1)` $b := y$

$(a, b) = (1, 1)$ ✓ $(a, b) = (1, 2)$ ✗

In this example, $(a = 1)$ means that $(z := x^2)$ and $(y := z^1)$ executed before $(a := y^2)$. Since $(z := x^2)$ read from the future operation $(x^2 := 1)$, necessarily $(a := y^2)$ has to consider $(y^2 := 2)$. $(a = 1)$ thus implies that $(y^2 := 2)$ executed before $(y := z^1)$, and that the final value of y is 1.

A.10 coreRMA comparison

coreRMA from Dan et al. [2016] does not allow for behaviours **SB3** and **2+2W3**, even without their In-Order Routing (IOR) axiom. These two examples rely on the fact that poll operations do not wait for $\mathbf{wb_R}$ to be empty. coreRMA also does not allow for behaviours such as **MP4**, where remote reads are reordered.

On the other hand, unlike our semantics, coreRMA allows (even with IOR) for the following three behaviours of Fig. 16, as discussed in Section 6.

RRPR3	
$x, y = 1$	$z = 0$
$x := z^2$ $y := z^2$ $a := y$ $b := x$	

$(a, b) = (0, 1)$ ✗

CRMA1		
$x = 0$	$y = 0$	$z = 0$
$x := y^2$ `flush(2)` $z^3 := x$ $a := x$	$y := 2$	$z := 1$ $b := z$

$(a, b) = (2, 0)$ ✗

CRMA2	
$x, y = 0$	$z, w = 0$
$z^2 := x$ $w^2 := y$ $y := 2$ $x := 2$	$w := 1$ $a := z$ $b := w$

$(a, b) = (2, 0)$ ✗

Fig. 16. coreRMA Valid Litmus Tests.

In **RRPR3**, coreRMA allows the local writes on x and y to be reordered, while in our semantics they must occur in order.

In **CRMA1**, coreRMA allows $(z^3 := x)$ to be performed before $(x := y^2)$, while in our semantics (if we replace `flush(2)` with `poll(2)`) we can stop the reordering.

In **CRMA2**, coreRMA allows the local read parts of $(z^2 := x)$ and $(w^2 := y)$ to be reordered, while RDMA$^{\text{TSO}}$ does not.

A.11 Observing pending writes in $\mathbf{wb_L}$ and $\mathbf{wb_R}$

NIC reads flush their corresponding $\mathbf{wb_L}$ or $\mathbf{wb_R}$ buffer. Intuitively, one could think that this would be enough to stop weak behaviours related to these buffers. But it is possible to observe events in-between the moment a write is sent to these buffers and before it is committed to memory.

GFP1	
$x = 0$	$y, z = 0$
$x := z^2$ `rfence` (2) $t := 1$ $y^2 := t$ $b := x$ $c := x$	$a := y$ $x^1 := 2$

$(a, b, c) = (1, 2, 0)$ ✗

GFP2	
$x = 0$	$y, z = 0$
$x := z^2$ `rfence` (2) $y^2 := 1$ $b := x$ $c := x$	$a := y$ $x^1 := 2$

$(a, b, c) = (1, 2, 0)$ ✓

GFG	
$x, y = 0$	$w, z = 0$
$x := z^2$ `rfence` (2) $y := w^2$ $a := x$ $b := x$ $c := y$	$w := 1$ $w := 2$ $x^1 := 3$

$(a, b, c) = (3, 0, 1)$ ✓*

Fig. 17. Observing $\mathbf{wb_R}$ and $\mathbf{wb_L}$ Litmus Tests.

As explained in the overview—notably with the example of Fig. 3d—a poll operation does not flush the $\mathbf{wb_R}$ buffer, which explain why modelling it is required. But the `rFence` operation does not flush the $\mathbf{wb_L}$ buffer either. The examples of Fig. 17 show some possible weak behaviours, explaining the need to model $\mathbf{wb_L}$.

The behaviour **GFP1** is not allowed, because the NIC local read of $(y^2 := t)$ flushes the write on z.

However, if we assume that $(y^2 := 1)$ does not perform a NIC read (i.e., inlining the data, which is allowed but not part of the definitions of this paper), then **GFP2** shows that, even with the remote fence, the later operation can fully execute before $(x := 0)$ is committed to memory. Note that this is not allowed by the semantics of this paper, as we do not model inlined data, and simplified the semantics by assuming puts necessarily perform a NIC local read.

Similarly, **GFG** shows that the remote read on w^2 can still happen before $(x := 0)$ is committed to memory. While allowed in theory, we did not manage to exhibit this behaviour in our testing.

A.12 Propensity of weak behaviours

Most of the weak behaviours presented in this appendix do not occur unless the whole system is stressed in specific ways. Table 18 summarises the techniques we employed for each test and how often we observe the weak behaviours once the setup has been adapted for this specific test.

Litmus test	Rate	Notes and techniques used
ST2	∼ 02%	No delay between the two operations
ST4	∼ 68%	Requires stress load (`ib_write_bw`)
ST5	∼ 02%	No need for stress load
ST6	∼ 40%	Had to stress the Completion Queues of both nodes
ST8	∼ 65%	Need random delay between the two operations
MP3	0%	This behaviour does not occur on current hardware.
MP3bis	∼ 100%	Systematic with background load on the system
MP4	0%	Have not observed it yet...
SB3	0%	This behaviour does not occur on current hardware.
SB3bis	∼ 100%	Systematic with stress load (`ib_write_bw`)
LB2	∼ 03%	Observed with random CPU delays combined with background load
LB3	∼ 98%	Very frequent with stress load (`ib_write_bw`)
IRIW1	∼ 05%	Occurs sometimes even without background load
2+2W3	0%	This behaviour does not occur on current hardware.
2+2W3bis	∼ 90%	Requires background load
QP1	0%	This behaviour does not occur on current hardware.
RFF1	< 0.5%	Occurs (rarely) even without background load
GFP2	rare	Observed on only one of our two setups, using 128 Queue Pairs
GFG	0%	Have not observed it yet...

Fig. 18. Occurrence rates of weak behaviours and testing notes

B SIMPLIFIED OPERATIONAL SEMANTICS

As presented in §3.2, all six buffers of the queue-pair pipeline are FIFO (entries are processed when they reach the head of a buffer) and they simply forward entries in order. While this structure closely mimics the hardware implementation, conceptually it is rather elaborate; indeed, equivalent semantics can be achieved with a simpler pipeline. In this section, we show an equivalent simplified operational semantics.

B.1 Definitions

$$\frac{\mathbf{sqp}.\mathbf{pipe} = \alpha \cdot (\mathtt{rfence}\ \overline{n})}{\mathsf{M}, \mathbf{sqp} \rightarrow_{\mathsf{sqp}} \mathsf{M}, \mathbf{sqp}[\mathbf{pipe} \mapsto \alpha]} \qquad \frac{\begin{array}{c}\mathbf{sqp}.\mathbf{pipe} = \alpha \cdot (\overline{y} := x) \cdot \beta \qquad \mathbf{wb_L} \in \{\mathsf{cn}\}^* \\ \beta \in \{\overline{y'} := v', x' := \overline{y'}, x' := v', \mathsf{ack_p}\}^*\end{array}}{\mathsf{M}, \mathbf{sqp} \rightarrow_{\mathsf{sqp}} \mathsf{M}, \mathbf{sqp}[\mathbf{pipe} \mapsto \alpha \cdot (\overline{y} := \mathsf{M}(x)) \cdot \beta]}$$

$$\frac{\begin{array}{c}\mathbf{sqp}.\mathbf{pipe} = \alpha \cdot (\overline{y} := v) \cdot \beta \qquad \beta \in \{x' := \overline{y'}, x' := v', \mathsf{ack_p}\}^* \\ \mathbf{sqp}' = \mathbf{sqp}[\mathbf{pipe} \mapsto \alpha \cdot \mathsf{ack_p} \cdot \beta][\mathbf{wb_R} \mapsto (\overline{y} := v) \cdot \mathbf{sqp}.\mathbf{wb_R}]\end{array}}{\mathsf{M}, \mathbf{sqp} \rightarrow_{\mathsf{sqp}} \mathsf{M}, \mathbf{sqp}'} \qquad \frac{\mathbf{sqp}.\mathbf{wb_R} = \alpha \cdot (\overline{y} := v)}{\mathsf{M}, \mathbf{sqp} \rightarrow_{\mathsf{sqp}} \mathsf{M}[\overline{y} \mapsto v], \mathbf{sqp}[\mathbf{wb_R} \mapsto \alpha]}$$

$$\frac{\begin{array}{c}\mathbf{sqp}.\mathbf{pipe} = \alpha \cdot \mathsf{ack_p} \\ \mathbf{sqp}' = \mathbf{sqp}[\mathbf{pipe} \mapsto \alpha][\mathbf{wb_L} \mapsto \mathsf{cn} \cdot \mathbf{sqp}.\mathbf{wb_L}]\end{array}}{\mathsf{M}, \mathbf{sqp} \rightarrow_{\mathsf{sqp}} \mathsf{M}, \mathbf{sqp}'} \qquad \frac{\begin{array}{c}\mathbf{sqp}.\mathbf{pipe} = \alpha \cdot (x := \overline{y}) \cdot \beta \qquad \beta \in \{x' := \overline{y'}, x' := v', \mathsf{ack_p}\}^* \\ \mathbf{sqp}.\mathbf{wb_R} = \varepsilon \qquad \mathbf{sqp}' = \mathbf{sqp}[\mathbf{pipe} \mapsto \alpha \cdot (x := \mathsf{M}(\overline{y})) \cdot \beta]\end{array}}{\mathsf{M}, \mathbf{sqp} \rightarrow_{\mathsf{sqp}} \mathsf{M}, \mathbf{sqp}'}$$

$$\frac{\begin{array}{c}\mathbf{sqp}.\mathbf{pipe} = \alpha \cdot (x := v) \\ \mathbf{sqp}' = \mathbf{sqp}[\mathbf{pipe} \mapsto \alpha][\mathbf{wb_L} \mapsto \mathsf{cn} \cdot (x := v) \cdot \mathbf{sqp}.\mathbf{wb_L}]\end{array}}{\mathsf{M}, \mathbf{sqp} \rightarrow_{\mathsf{sqp}} \mathsf{M}, \mathbf{sqp}'} \qquad \frac{\begin{array}{c}\mathbf{sqp}.\mathbf{wb_L} = \alpha \cdot (x := v) \cdot \beta \quad \beta \in \{\mathsf{cn}\}^* \\ \mathbf{sqp}' = \mathbf{sqp}[\mathbf{wb_L} \mapsto \alpha \cdot \beta]\end{array}}{\mathsf{M}, \mathbf{sqp} \rightarrow_{\mathsf{sqp}} \mathsf{M}[x \mapsto v], \mathbf{sqp}'}$$

Fig. 19. Queue-pair transitions of the simplified RDMA$^{\mathsf{TSO}}$ operational semantics

We replace the four buffers in the middle of a queue pair, namely $\mathbf{req_L}$, $\mathbf{in_R}$, $\mathbf{out_R}$ and $\mathbf{rsp_L}$, with a *single buffer*, **pipe**, resulting in a queue pair comprising **pipe**, $\mathbf{wb_L}$ and $\mathbf{wb_R}$. We retain $\mathbf{wb_L}$ and $\mathbf{wb_R}$ as separate entities to model the delayed propagation of NIC local (in $\mathbf{wb_L}$) and remote (in $\mathbf{wb_R}$) writes.

We note $\mathsf{Pipe}_n^{\overline{n}} \triangleq \{y^{\overline{n}} := x^n, y^{\overline{n}} := v, \mathsf{ack_p}, x^n := y^{\overline{n}}, x^n := v, \mathtt{rfence}\ \overline{n}\}^*$ the type of pipe buffers from a node n towards a remote note $\overline{n}$. A pipe buffer can contain puts, gets, simplified puts and gets (with values), acknowledgements, and remote fences.

A *simple queue pair* is then modeled as a tuple $\mathbf{sqp} \in \mathsf{SQPair}_n^{\overline{n}} \triangleq \mathsf{Pipe}_n^{\overline{n}} \times \mathsf{WBR}_n^{\overline{n}} \times \mathsf{WBL}_n^{\overline{n}}$.

Given a sequence α and a set S, we write $\alpha \in S^*$ to denote that all entries in α are drawn from S. The simplified queue-pair transitions in Fig. 19 operate on simplified queue pairs and describe the progress of remote operations through the simplified pipeline. Specifically, as $\mathbf{in_R}$, $\mathbf{out_R}$ and $\mathbf{rsp_L}$ are now captured by **pipe**, we process an `rFence` by simply removing it when it is at the head of **pipe** (first transition). The next four transitions describe how a put $\overline{y} := x$ is processed by obtaining the value of x (second transition, capturing step P_1), forwarding it to $\mathbf{wb_R}$ (third transition, step P_2), propagating it to the memory (fourth transition, steps P_3–P_4), and producing its completion notification (fifth transition, step P_5). Analogously, the next three transitions describe how a get $x := \overline{y}$ is processed by fulfilling the value of $\overline{y}$ (sixth transition, capturing steps G_1–G_4), forwarding it to $\mathbf{wb_L}$ (penultimate transition, step G_5) and propagating its result to memory (last transition, step G_6).

$$\mathsf{M} \in \mathsf{Mem} \triangleq \mathsf{Loc} \to \mathsf{Val} \qquad \mathsf{B} \in \mathsf{SBMap} \triangleq \lambda t \in \mathsf{Tid}.\mathsf{SBuff}_{n(t)}$$
$$\mathsf{QP} \in \mathsf{SQPMap} \triangleq \lambda t.\big(\lambda \overline{n(t)}.\mathsf{SQPair}_n^{\overline{n}}\big) \qquad \mathsf{b} \in \mathsf{SBuff}_n \triangleq \big\{x^n{:=}v, y^{\overline{n}}{:=}x^n, x^n{:=}y^{\overline{n}}, \mathtt{rfence}\ \overline{n}\big\}^*$$
$$\mathbf{sqp} \in \mathsf{SQPair}_n^{\overline{n}} \triangleq \mathsf{Pipe}_n^{\overline{n}} \times \mathsf{WBR}_n^{\overline{n}} \times \mathsf{WBL}_n^{\overline{n}} \qquad \mathbf{wb_R} \in \mathsf{WBR}_n^{\overline{n}} \triangleq \big\{y^{\overline{n}} := v\big\}^*$$
$$\mathbf{pipe} \in \mathsf{Pipe}_n^{\overline{n}} \triangleq \big\{y^{\overline{n}} := x^n, y^{\overline{n}} := v, \mathsf{ack_p}, x^n := y^{\overline{n}}, x^n := v, \mathtt{rfence}\ \overline{n}\big\}^* \qquad \mathbf{wb_L} \in \mathsf{WBL}_n^{\overline{n}} \triangleq \big\{\mathsf{cn}, x^n := v\big\}^*$$

$$\frac{\mathsf{B}' = \mathsf{B}[t \mapsto (x := v) \cdot \mathsf{B}(t)]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\mathtt{lW}(x,v)} \mathsf{M}, \mathsf{B}', \mathsf{QP}} \qquad \frac{(\mathsf{M} \lhd \mathsf{B}(t))(x) = v}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\mathtt{lR}(x,v)} \mathsf{M}, \mathsf{B}, \mathsf{QP}} \qquad \frac{\mathsf{B}(t) = \varepsilon \qquad \mathsf{M}(x) = v_1}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\mathsf{CASS}(x,v_1,v_2)} \mathsf{M}[x \mapsto v_2], \mathsf{B}, \mathsf{QP}}$$

$$\frac{\mathsf{B}(t) = \varepsilon \qquad \mathsf{M}(x) = v}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\mathsf{CASF}(x,v)} \mathsf{M}, \mathsf{B}, \mathsf{QP}} \qquad \frac{\mathsf{B}(t) = \varepsilon}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\mathsf{F}} \mathsf{M}, \mathsf{B}, \mathsf{QP}} \qquad \frac{\mathsf{B}(t) = \mathsf{b} \cdot (x := v)}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\varepsilon} \mathsf{M}[x \mapsto v], \mathsf{B}[t \mapsto \mathsf{b}], \mathsf{QP}}$$

$$\frac{\mathsf{rc}^n \in \big\{x := y^{\overline{n}}, y^{\overline{n}} := x, \mathtt{rfence}\ \overline{n}\big\} \qquad \begin{array}{c}\mathsf{B}(t) = \mathsf{b} \cdot \mathsf{rc}^n \\ \mathsf{QP}(t)(\overline{n}) = \mathbf{sqp}\end{array} \qquad \mathbf{sqp}' = \mathbf{sqp}[\mathbf{pipe} \mapsto \mathsf{rc}^n \cdot \mathbf{sqp}.\mathbf{pipe}]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\varepsilon} \mathsf{M}, \mathsf{B}[t \mapsto \mathsf{b}], \mathsf{QP}[t \mapsto \mathsf{QP}(t)[\overline{n} \mapsto \mathbf{sqp}']]}$$

$$\frac{\mathsf{B}' = \mathsf{B}[t \mapsto (x := \overline{y}) \cdot \mathsf{B}(t)]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\mathsf{Get}(x,\overline{y})} \mathsf{M}, \mathsf{B}', \mathsf{QP}} \qquad \frac{\mathsf{B}' = \mathsf{B}[t \mapsto (\overline{y} := x) \cdot \mathsf{B}(t)]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\mathsf{Put}(\overline{y},x)} \mathsf{M}, \mathsf{B}', \mathsf{QP}} \qquad \frac{\mathsf{B}' = \mathsf{B}[t \mapsto (\mathtt{rfence}\ \overline{n}) \cdot \mathsf{B}(t)]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\mathsf{rF}(\overline{n})} \mathsf{M}, \mathsf{B}', \mathsf{QP}}$$

$$\frac{\mathsf{QP}(t)(\overline{n}) = \mathbf{sqp} \qquad \mathbf{sqp}.\mathbf{wb_L} = \alpha \cdot \mathsf{cn} \qquad \mathbf{sqp}' = \mathbf{sqp}[\mathbf{wb_L} \mapsto \alpha]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\mathsf{P}(\overline{n})} \mathsf{M}, \mathsf{B}, \mathsf{QP}[t \mapsto \mathsf{QP}(t)[\overline{n} \mapsto \mathbf{sqp}']]} \qquad \frac{\mathsf{M}, \mathsf{QP}(t)(\overline{n}) \to_{\mathsf{sqp}} \mathsf{M}', \mathbf{sqp} \quad \text{(Fig. 19)}}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{t:\varepsilon} \mathsf{M}', \mathsf{B}, \mathsf{QP}[t \mapsto \mathsf{QP}(t)[\overline{n} \mapsto \mathbf{sqp}]]}$$

$$\text{with} \quad (\mathsf{M} \lhd \alpha)(x) \triangleq \begin{cases} v & \text{if } \alpha = \beta \cdot (x := v) \cdot - \land \forall v'.\, x := v' \notin \beta \\ \mathsf{M}(x) & \text{if } \forall v.\, x := v \notin \alpha \end{cases}$$

Fig. 20. RDMA$^{\mathsf{TSO}}$ simplified hardware domains (above) and hardware transitions (below)

We keep the program transitions from the concrete operational semantics (Fig. 6). The hardware transitions of our simplified semantics, in Fig. 20, is almost identical to the previous one (Fig. 7). The only two differences are: (1) we use the simplified queue-pair transitions from Fig. 19; (2) new remote operations are added to the beginning of **pipe** (middle rule), since $\mathbf{req_L}$ no longer exists.

The simplified operational semantics then simply combines the program and hardware transition, similarly to Fig. 8.

B.2 Equivalence Proof

The simplified semantics described above, using the new queue-pair transitions (Fig. 19), is equivalent to the semantics defined in Section 3.2. The proof has been formalised in the Coq proof assistant. In this section, we provide an overview of the equivalence proof.

We split the previous queue-pair semantics (from Fig. 7) into two sets of rules in Fig. 21. ($\to_{\mathsf{s}}$) is used when an action is moved from the end of a buffer to the beginning of the next. ($\to_{\mathsf{u}}$) is used for visible actions: reads, writes, and actions of the $\mathbf{wb_R}$ and $\mathbf{wb_L}$ buffers. Since shifting rules do not modify the memory, we simplify them into a relation from queue pair to queue pair.

Definition B.1. We say that a 6-buffers queue pair $\langle \mathbf{req_L}, \mathbf{in_R}, \mathbf{wb_R}, \mathbf{out_R}, \mathbf{rsp_L}, \mathbf{wb_L} \rangle$ is well-formed if the four main buffers respect the conditions of Figure 5. I.e.:

- $\mathbf{req_L} \in \big\{\overline{y} := x, x := \overline{y}, \mathtt{rfence}\ \overline{n}\big\}^*$
- $\mathbf{in_R} \in \big\{\overline{y} := v, x := \overline{y}\big\}^*$

Internal Shifting Steps

$$\dfrac{\mathbf{req}_{\mathrm{L}} = \mathbf{req}'_{\mathrm{L}} \cdot (x := \overline{y}) \qquad \mathbf{in}'_{\mathrm{R}} = (x := \overline{y}) \cdot \mathbf{in}_{\mathrm{R}}}{\langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{s}} \langle \mathbf{req}'_{\mathrm{L}}, \mathbf{in}'_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{in}_{\mathrm{R}} = \mathbf{in}'_{\mathrm{R}} \cdot (x := \overline{y}) \qquad \mathbf{out}'_{\mathrm{R}} = (x := \overline{y}) \cdot \mathbf{out}_{\mathrm{R}}}{\langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{s}} \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}'_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}'_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{out}_{\mathrm{R}} = \mathbf{out}'_{\mathrm{R}} \cdot (x := v) \qquad \mathbf{rsp}'_{\mathrm{L}} = (x := v) \cdot \mathbf{rsp}_{\mathrm{L}}}{\langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{s}} \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}'_{\mathrm{R}}, \mathbf{rsp}'_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{out}_{\mathrm{R}} = \mathbf{out}'_{\mathrm{R}} \cdot (\mathsf{ack}_{\mathsf{p}}) \qquad \mathbf{rsp}'_{\mathrm{L}} = (\mathsf{ack}_{\mathsf{p}}) \cdot \mathbf{rsp}_{\mathrm{L}}}{\langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{s}} \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}'_{\mathrm{R}}, \mathbf{rsp}'_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle}$$

Visible Steps

$$\dfrac{\mathbf{req}_{\mathrm{L}} = \mathbf{req}'_{\mathrm{L}} \cdot (\mathtt{rfence}\ \overline{n}) \qquad \mathbf{in}_{\mathrm{R}} = \mathbf{out}_{\mathrm{R}} = \mathbf{rsp}_{\mathrm{L}} = \varepsilon}{\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{u}} \mathsf{M}, \langle \mathbf{req}'_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{req}_{\mathrm{L}} = \mathbf{req}'_{\mathrm{L}} \cdot (\overline{y} := x) \qquad \mathbf{in}'_{\mathrm{R}} = (\overline{y} := \mathsf{M}(x)) \cdot \mathbf{in}_{\mathrm{R}} \qquad \mathbf{wb}_{\mathrm{L}} \in \left\{\mathsf{cn}\right\}^*}{\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{u}} \mathsf{M}, \langle \mathbf{req}'_{\mathrm{L}}, \mathbf{in}'_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{in}_{\mathrm{R}} = \mathbf{in}'_{\mathrm{R}} \cdot (\overline{y} := v) \qquad \mathbf{wb}'_{\mathrm{R}} = (\overline{y} := v) \cdot \mathbf{wb}_{\mathrm{R}} \qquad \mathbf{out}'_{\mathrm{R}} = (\mathsf{ack}_{\mathsf{p}}) \cdot \mathbf{out}_{\mathrm{R}}}{\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{u}} \mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}'_{\mathrm{R}}, \mathbf{wb}'_{\mathrm{R}}, \mathbf{out}'_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{wb}_{\mathrm{R}} = \mathbf{wb}'_{\mathrm{R}} \cdot (\overline{y} := v)}{\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{u}} \mathsf{M}[\overline{y} \mapsto v], \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}'_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{rsp}_{\mathrm{L}} = \mathbf{rsp}'_{\mathrm{L}} \cdot (\mathsf{ack}_{\mathsf{p}}) \qquad \mathbf{wb}'_{\mathrm{L}} = \mathsf{cn} \cdot \mathbf{wb}_{\mathrm{L}}}{\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{u}} \mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}'_{\mathrm{L}}, \mathbf{wb}'_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{out}_{\mathrm{R}} = \alpha \cdot (x := \overline{y}) \cdot \beta \qquad \mathbf{wb}_{\mathrm{R}} = \varepsilon \qquad \mathbf{out}'_{\mathrm{R}} = \alpha \cdot (x := \mathsf{M}(\overline{y})) \cdot \beta}{\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{u}} \mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}'_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{rsp}_{\mathrm{L}} = \mathbf{rsp}'_{\mathrm{L}} \cdot (x := v) \qquad \mathbf{wb}'_{\mathrm{L}} = \mathsf{cn} \cdot (x := v) \cdot \mathbf{wb}_{\mathrm{L}}}{\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{u}} \mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}'_{\mathrm{L}}, \mathbf{wb}'_{\mathrm{L}} \rangle}$$

$$\dfrac{\mathbf{wb}_{\mathrm{L}} = \alpha \cdot (x := v) \cdot \beta \qquad \beta \in \left\{\mathsf{cn}\right\}^* \qquad \mathbf{wb}'_{\mathrm{L}} = \alpha \cdot \beta}{\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \longrightarrow_{\mathsf{u}} \mathsf{M}[x \mapsto v], \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}'_{\mathrm{L}} \rangle}$$

Fig. 21. 6 Buffers NIC Semantics

- $\mathbf{out}_{\mathrm{R}} \in \left\{x := \overline{y}, x := v, \mathsf{ack}_{\mathsf{p}}\right\}^*$
- $\mathbf{rsp}_{\mathrm{L}} \in \left\{x := v, \mathsf{ack}_{\mathsf{p}}\right\}^*$

We also say that a queue-pair map (and other structures) is well-formed if it only contains well-formed queue pairs.

Lemma B.2. *Empty queue pairs (maps mapping only to empty queue pairs) are well-formed.*

Proof. trivial $\square$

Lemma B.3. *The different relations presented in this document preserve well-formedness. E.g., if* $\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP} \Rightarrow \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}'$ *and* QP *is well-formed, then* QP' *is well-formed.*

Proof. By cases analysis on every induction rule. $\square$

Definition B.4. We note $[\![\mathsf{qp}]\!]$ the merging of the 6-buffers queue pair into its 3-buffers counterpart. I.e., $[\![\langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle]\!] \triangleq \langle \mathbf{req}_{\mathrm{L}} \cdot \mathbf{in}_{\mathrm{R}} \cdot \mathbf{out}_{\mathrm{R}} \cdot \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{L}} \rangle$.

We reuse the notation for maps ($[\![\mathsf{QP}]\!]$) when merging all their queue pairs at the same time.

Definition B.5. We note $\mathbf{unmerge}(\mathbf{sqp})$ the reverse operation, turning a 3-buffer queue pair into its 6-buffers counterpart. $\mathbf{wb}_{\mathrm{R}}$ and $\mathbf{wb}_{\mathrm{L}}$ are left unchanged. The $\mathbf{pipe}$ buffer is split into $(\mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}})$. To proceed, we consider each element of $\mathbf{pipe}$ one-by-one in reverse order (right to left). We put as many elements as possible in $\mathbf{rsp}_{\mathrm{L}}$. as soon as one element e cannot be put in $\mathbf{rsp}_{\mathrm{L}}$ (i.e., $e \notin \{x := v, \mathsf{ack}_{\mathsf{p}}\}$), we start putting as many elements as possible in $\mathbf{out}_{\mathrm{R}}$. Once again, when we cannot ($e \notin \{x := \overline{y}, x := v, \mathsf{ack}_{\mathsf{p}}\}$) we start filling $\mathbf{in}_{\mathrm{R}}$. And when we cannot ($e \notin \{\overline{y} := v, x := \overline{y}\}$), we put the remaining elements in $\mathbf{req}_{\mathrm{L}}$.

Lemma B.6. *For all 3-buffer queue pairs* $\mathbf{sqp}$*,* $[\![\mathbf{unmerge}(\mathbf{sqp})]\!] = \mathbf{sqp}$.

Proof. trivial $\square$

Lemma B.7. *For all queue pairs* qp *and* qp'*, if* $\mathsf{qp} \rightarrow_{\mathsf{s}} \mathsf{qp}'$ *then* $[\![\mathsf{qp}]\!] = [\![\mathsf{qp}']\!]$.

Proof. trivial $\square$

Lemma B.8. *For all memories* M *and* M' *and well-formed queue pairs* qp *and* qp'*,*
if $\mathsf{M}, \mathsf{qp} \rightarrow_{\mathsf{u}} \mathsf{M}', \mathsf{qp}'$ *then* $\mathsf{M}, [\![\mathsf{qp}]\!] \rightarrow_{\mathsf{sqp}} \mathsf{M}', [\![\mathsf{qp}']\!]$.

Proof. For each possible rule of Figure 21, we need to show that the conditions are met to apply the corresponding rule from Figure 19.

- If $\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}} \cdot (\mathtt{rfence}\ \overline{n}), \varepsilon, \mathbf{wb}_{\mathrm{R}}, \varepsilon, \varepsilon, \mathbf{wb}_{\mathrm{L}} \rangle \rightarrow_{\mathsf{u}} \mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \varepsilon, \mathbf{wb}_{\mathrm{R}}, \varepsilon, \varepsilon, \mathbf{wb}_{\mathrm{L}} \rangle$, then $[\![\mathsf{qp}]\!] = \langle \mathbf{req}_{\mathrm{L}} \cdot (\mathtt{rfence}\ \overline{n}), \mathbf{wb}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{L}} \rangle$ and $[\![\mathsf{qp}']\!] = \langle \mathbf{req}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{L}} \rangle$. We can apply the first rule of Figure 19 to derive $\mathsf{M}, [\![\mathsf{qp}]\!] \rightarrow_{\mathsf{sqp}} \mathsf{M}, [\![\mathsf{qp}']\!]$
- If $\mathsf{M}, \langle \alpha \cdot (\overline{y} := x), \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \varepsilon \rangle \rightarrow_{\mathsf{u}} \mathsf{M}, \langle \alpha, (\overline{y} := \mathsf{M}(x)) \cdot \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \varepsilon \rangle$, then we can apply the second rule. The condition to check is $\mathbf{in}_{\mathrm{R}} \cdot \mathbf{out}_{\mathrm{R}} \cdot \mathbf{rsp}_{\mathrm{L}} \in \left\{\overline{y} := v, x := \overline{y}, x := v, \mathsf{ack}_{\mathsf{p}}\right\}^*$, which comes from Definition B.1.
- Similarly to case 2
- If $\mathsf{M}, \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \alpha \cdot (\overline{y} := v), \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle \rightarrow_{\mathsf{u}} \mathsf{M}[\overline{y} \mapsto v], \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \alpha, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle$, then we similarly have $\mathsf{M}, \langle \mathbf{pipe}, \alpha \cdot (\overline{y} := v), \mathbf{wb}_{\mathrm{L}} \rangle \rightarrow_{\mathsf{sqp}} \mathsf{M}[\overline{y} \mapsto v], \langle \mathbf{pipe}, \alpha, \mathbf{wb}_{\mathrm{L}} \rangle$ with the fourth rule of Figure 19, with $\mathbf{pipe} = \mathbf{req}_{\mathrm{L}} \cdot \mathbf{in}_{\mathrm{R}} \cdot \mathbf{out}_{\mathrm{R}} \cdot \mathbf{rsp}_{\mathrm{L}}$.
- Similarly to case 2
- Similarly to case 2
- Similarly to case 4

$\square$

Theorem B.9. *For all well-formed queue-pair map* QP*, if* $\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP} \Rightarrow \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}'$*, then either* $\mathsf{P} = \mathsf{P}'$*,* $\mathsf{M} = \mathsf{M}'$*,* $\mathsf{B} = \mathsf{B}'$*, and* $[\![\mathsf{QP}]\!] = [\![\mathsf{QP}']\!]$ *(zero steps), or* $\mathsf{P}, \mathsf{M}, \mathsf{B}, [\![\mathsf{QP}]\!] \Rightarrow \mathsf{P}', \mathsf{M}', \mathsf{B}', [\![\mathsf{QP}']\!]$ *(one step).*

Proof. If it comes from an NIC reduction $\rightarrow_{\mathsf{s}}$, we use Lemma B.7. If it comes from an NIC reduction $\rightarrow_{\mathsf{u}}$, we use Lemma B.8. Otherwise, other reduction rules are similar in the simplified semantics and can be mapped one to one. □

Theorem B.9 above show that the simplified 3-buffers semantics preserves the behaviours of the previous 6-buffers semantics. Now, we need to show the converse.

Lemma B.10. *For all well-formed queue pairs* qp, $\mathsf{qp} \rightarrow_{\mathsf{s}}^{*} \mathbf{unmerge}([\![\mathsf{qp}]\!])$.

Proof. let $\mathsf{qp} = \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle$ and $\mathbf{unmerge}([\![\mathsf{qp}]\!]) = \langle \mathbf{req}'_{\mathrm{L}}, \mathbf{in}'_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}'_{\mathrm{R}}, \mathbf{rsp}'_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle$. We proceed by successive induction on (the last element of) $\mathbf{rsp}_{\mathrm{L}}$, $\mathbf{out}_{\mathrm{R}}$, $\mathbf{in}_{\mathrm{R}}$, and $\mathbf{req}_{\mathrm{L}}$. Unmerging pushes elements as far down the queue pair as possible. As soon as there is a mismatch, it means that unmerging places the element further, e.g. in $\mathbf{rsp}'_{\mathrm{L}}$ while the element comes from $\mathbf{out}_{\mathrm{R}}$. In each case, we check that we have the appropriate inference rule to shift this element. □

Lemma B.11. *For all well-formed queue pair* qp, *3-buffers queue pair* $\mathbf{sqp}'$, *memories* M *and* M', *if* $\mathsf{M}, [\![\mathsf{qp}]\!] \rightarrow_{\mathsf{sqp}} \mathsf{M}', \mathbf{sqp}'$ *then there exists* qp' *such that* $[\![\mathsf{qp}']\!] = \mathbf{sqp}'$ *and* $\mathsf{M}, \mathbf{unmerge}([\![\mathsf{qp}]\!]) \rightarrow_{\mathsf{u}} \mathsf{M}', \mathsf{qp}'$.

Proof. Let us note $\mathbf{unmerge}([\![\mathsf{qp}]\!]) = \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle$ and $\mathbf{sqp}' = \langle \mathbf{pipe}, \mathbf{wb}'_{\mathrm{R}}, \mathbf{wb}'_{\mathrm{L}} \rangle$. Note that $[\![\mathsf{qp}]\!] = [\![\mathbf{unmerge}([\![\mathsf{qp}]\!])]\!]$ by Lemma B.6.

- If $\mathsf{M}, [\![\mathsf{qp}]\!] \rightarrow_{\mathsf{sqp}} \mathsf{M}', \mathbf{sqp}'$ comes from the first rule, then $\mathbf{req}_{\mathrm{L}} \cdot \mathbf{in}_{\mathrm{R}} \cdot \mathbf{out}_{\mathrm{R}} \cdot \mathbf{rsp}_{\mathrm{L}}$ is of the form $\alpha \cdot (\texttt{rfence}\ \overline{n})$, and $\mathbf{pipe} = \alpha$, $\mathbf{wb}_{\mathrm{R}} = \mathbf{wb}'_{\mathrm{R}}$, and $\mathbf{wb}_{\mathrm{L}} = \mathbf{wb}'_{\mathrm{L}}$. By Definition B.5 for unmerging, $\mathbf{req}_{\mathrm{L}} = \alpha \cdot (\texttt{rfence}\ \overline{n})$, and $\mathbf{in}_{\mathrm{R}} = \mathbf{out}_{\mathrm{R}} = \mathbf{rsp}_{\mathrm{L}} = \varepsilon$. We use $\mathsf{qp}' = \langle \alpha, \varepsilon, \mathbf{wb}_{\mathrm{R}}, \varepsilon, \varepsilon, \mathbf{wb}_{\mathrm{L}} \rangle$. $[\![\mathsf{qp}']\!] = \mathbf{sqp}'$ holds, and we have $\mathsf{M}, \mathsf{qp} \rightarrow_{\mathsf{u}} \mathsf{M}', \mathsf{qp}'$ using the rule for $\texttt{rfence}\ \overline{n}$.
- If $\mathsf{M}, [\![\mathsf{qp}]\!] \rightarrow_{\mathsf{sqp}} \mathsf{M}', \mathbf{sqp}'$ comes from the second rule, then $\mathbf{req}_{\mathrm{L}} \cdot \mathbf{in}_{\mathrm{R}} \cdot \mathbf{out}_{\mathrm{R}} \cdot \mathbf{rsp}_{\mathrm{L}}$ is of the form $\alpha \cdot (\overline{y} := x) \cdot \beta$, with $\beta \in \left\{\overline{y'} := v', x' := \overline{y'}, x' := v', \mathsf{ack}_{\mathsf{p}}\right\}^{*}$, and $\mathbf{pipe} = \alpha \cdot (\overline{y} := \mathsf{M}(x)) \cdot \beta$, $\mathbf{wb}_{\mathrm{R}} = \mathbf{wb}'_{\mathrm{R}}$, and $\mathbf{wb}_{\mathrm{L}} = \mathbf{wb}'_{\mathrm{L}} \in \left\{\mathsf{cn}\right\}^{*}$. By Definition B.5 for unmerging, we necessarily have $\mathbf{req}_{\mathrm{L}} = \alpha \cdot (\overline{y} := x)$ and $\beta = \mathbf{in}_{\mathrm{R}} \cdot \mathbf{out}_{\mathrm{R}} \cdot \mathbf{rsp}_{\mathrm{L}}$. We use $\mathsf{qp}' = \langle \alpha, (\overline{y} := \mathsf{M}(x)) \cdot \mathbf{in}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{R}}, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle$. $[\![\mathsf{qp}']\!] = \mathbf{sqp}'$ holds, and we have $\mathsf{M}, \mathsf{qp} \rightarrow_{\mathsf{u}} \mathsf{M}', \mathsf{qp}'$ using the rule for $(\overline{y} := x)$.
- Similarly to case 2
- If $\mathsf{M}, [\![\mathsf{qp}]\!] \rightarrow_{\mathsf{sqp}} \mathit{Mem}', \mathbf{sqp}'$ comes from the fifth rule, then $\mathbf{wb}_{\mathrm{R}}$ is of the form $\alpha \cdot (\overline{y} := v)$, with $\mathsf{M}' = \mathsf{M}[\overline{y} \mapsto v]$ and $\mathbf{sqp}' = \mathbf{req}_{\mathrm{L}} \cdot \mathbf{in}_{\mathrm{R}} \cdot \mathbf{out}_{\mathrm{R}} \cdot \mathbf{rsp}_{\mathrm{L}}$ and $\mathbf{wb}'_{\mathrm{R}} = \alpha$. We simply choose $\mathsf{qp}' = \langle \mathbf{req}_{\mathrm{L}}, \mathbf{in}_{\mathrm{R}}, \alpha, \mathbf{out}_{\mathrm{R}}, \mathbf{rsp}_{\mathrm{L}}, \mathbf{wb}_{\mathrm{L}} \rangle$. $[\![\mathsf{qp}']\!] = \mathbf{sqp}'$ holds, and we have $\mathsf{M}, \mathsf{qp} \rightarrow_{\mathsf{u}} \mathsf{M}', \mathsf{qp}'$ using the rule for $\mathbf{wb}_{\mathrm{R}}$.
- Similarly to case 2
- Similarly to case 2
- Similarly to case 4

□

Theorem B.12. *For all well-formed queue-pair map* QP *(of the concrete semantics) and merged queue-pair map* QP'' *(of the simplified semantics), if* $\mathsf{P}, \mathsf{M}, \mathsf{B}, [\![\mathsf{QP}]\!] \Rightarrow \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}''$, *then there exists* QP' *such that* $[\![\mathsf{QP}']\!] = \mathsf{QP}''$ *and* $\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP} \Rightarrow^{*} \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}'$.

Proof. By case analysis on the rule used. If it comes from an NIC reduction $\rightarrow_{\mathsf{sqp}}$, we use both Lemmas B.10 and B.11. Otherwise, other cases of the simplified semantics can be matched by a similar step of the main semantics. □

All of the results above are formalised in the Coq proof assistant and are available in the Supplementary Material.

C APPENDIX TO THE DECLARATIVE SEMANTICS

C.1 Example Execution Graph

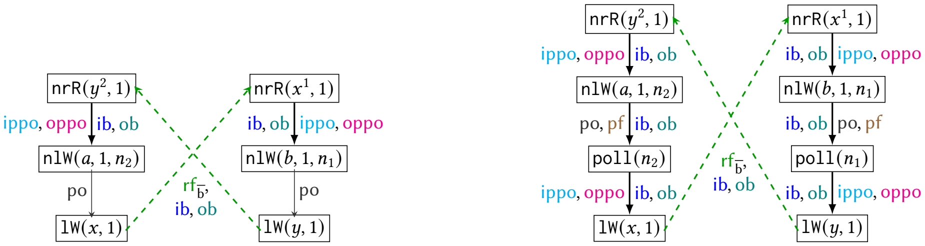

(a) An execution of Fig. 3b with outcome $a=b=1$

(b) An execution of Fig. 3c with outcome $a=b=1$

Fig. 22. Examples of consistent (a) and inconsistent (b) execution graphs

Example Executions. We depict two example executions in Fig. 22, where each event is represented as a rectangle with the corresponding label, and each column represents the events of one thread. Specifically, Fig. 22a shows the execution of Fig. 3b with outcome $a = b = 1$, and Fig. 22b shows the execution of Fig. 3c with outcome $a = b = 1$. Note that each get/put operation is represented as two events. Each edge is labelled with one or more relations; e.g. the top-left NIC remote read in Fig. 22a (labelled nrR, associated with the $a := y^2$ get) is related to the subsequent NIC local write (labelled nlW, also associated with $a := y^2$) via ippo, oppo, ib and ob. For brevity, we have omitted the po labels when they are also either ippo- or oppo-related (as oppo $\subseteq$ ippo $\subseteq$ po).

Note that both graphs in Fig. 22 are valid executions according to Def. 4.3. However, while Fig. 22a *is* a consistent execution, the one in Fig. 22b *is not*. Specifically, the execution in Fig. 22b is not consistent as both relations ib and ob contain a cycle.

C.2 Event Sequence Construction

$$\frac{C \rightsquigarrow C' \qquad C' \rightarrowtail s}{C \rightarrowtail s} \qquad \frac{C_1 \rightarrowtail s_1 \qquad C_2 \rightarrowtail s_2}{C_1; C_2 \rightarrowtail s_1, s_2} \qquad \frac{\mathrm{elocs}(e) = \emptyset}{x := e \rightarrowtail \mathtt{lW}(x, [\![e]\!])}$$

$$\frac{\mathrm{elocs}(e_{\mathrm{old}}) = \mathrm{elocs}(e_{\mathrm{new}}) = \emptyset \quad z := [\![e_{\mathrm{old}}]\!] \rightarrowtail s}{z := \mathtt{CAS}(x, e_{\mathrm{old}}, e_{\mathrm{new}}) \rightarrowtail \mathtt{CAS}(x, [\![e_{\mathrm{old}}]\!], [\![e_{\mathrm{new}}]\!]), s} \qquad \frac{\mathrm{elocs}(e_{\mathrm{old}}) = \mathrm{elocs}(e_{\mathrm{new}}) = \emptyset \quad v \neq [\![e_{\mathrm{old}}]\!] \qquad z := v \rightarrowtail s}{z := \mathtt{CAS}(x, e_{\mathrm{old}}, e_{\mathrm{new}}) \rightarrowtail \mathtt{F}, \mathtt{lR}(x, v), s}$$

$$\frac{}{\mathtt{mfence} \rightarrowtail \mathtt{F}} \qquad \frac{}{\mathtt{assume}(x = v) \rightarrowtail \mathtt{lR}(x, v)} \qquad \frac{v' \neq v}{\mathtt{assume}(x \neq v) \rightarrowtail \mathtt{lR}(x, v')}$$

$$\frac{}{x := y^{\overline{n}} \rightarrowtail \mathtt{nrR}(y^{\overline{n}}, v), \mathtt{nlW}(x, v, \overline{n})} \qquad \frac{}{y^{\overline{n}} := x \rightarrowtail \mathtt{nlR}(x, v, \overline{n}), \mathtt{nrW}(y^{\overline{n}}, v)}$$

$$\frac{}{\mathtt{rfence}\ (\overline{n}) \rightarrowtail \mathtt{nF}(\overline{n})} \qquad \frac{}{\mathtt{poll}(\overline{n}) \rightarrowtail \mathtt{P}(\overline{n})} \qquad \frac{}{\mathtt{skip} \rightarrowtail \epsilon}$$

Fig. 23. Label Sequences Construction

We define the following process to build and compose event graphs:

- For a thread identifier $t \in \mathsf{Tid}$ and a sequence of labels $l_1, \ldots, l_n \in \mathsf{ELab}$, we define $G^t(l_1, \ldots, l_n)$ as the set of all event graphs of the form $(\{\mathsf{e}_1, \ldots, \mathsf{e}_n\}, \mathsf{po})$ respecting the following conditions:
  - $l(\mathsf{e}_i) = l_i$ for every $1 \leq i \leq n$;
  - $\iota(\mathsf{e}_i) \neq \iota(\mathsf{e}_j)$ for every $1 \leq i < j \leq n$;
  - $t(\mathsf{e}_i) = t$ for every $1 \leq i \leq n$;
  - $\mathsf{po} = \left\{(\mathsf{e}_i, \mathsf{e}_j) \mid 1 \leq i < j \leq n\right\}$.
- The sequential composition of two event graphs G_1 and G_2, denoted $G_1; G_2$, is the event graph given by $(G_1.\mathsf{Event} \uplus G_2.\mathsf{Event}, G_1.\mathsf{po} \cup G_2.\mathsf{po} \cup G_1.\mathsf{Event} \times G_2.\mathsf{Event})$. Note that $G_1; G_2$ is only defined if $G_1.\mathsf{Event}$ and $G_2.\mathsf{Event}$ are disjoint, and we assume a disjoint set of event identifiers.
- The parallel composition of two event graphs G_1 and G_2, denoted $G_1 \parallel G_2$, is the event graph given by $(G_1.\mathsf{Event} \uplus G_2.\mathsf{Event}, G_1.\mathsf{po} \cup G_2.\mathsf{po})$. Note that $G_1 \parallel G_2$ is only defined if $G_1.\mathsf{Event}$ and $G_2.\mathsf{Event}$ are disjoint, and we assume a disjoint set of event identifiers. Parallel composition (which is commutative and associative) is generalised to apply on sets of event graphs (e.g., $\|_{i \in I}\, \mathcal{G}_i$) in the obvious way.

For a set Loc of locations, we say that $G_{\mathsf{init}} = (\mathsf{Event}_0, \emptyset)$ is an initial event graph if Event_0 is a set of initialisation events, with different event identifiers, with one label $\mathtt{lW}(x, 0)$ per location $x \in \mathsf{Loc}$.

We say that G is generated by a program P if $G = G_{\mathsf{init}}; (\|_{t \in \mathsf{Tid}}\, G_t)$ and there is sequences s_t such that: G_{init} is an initial event graph; $\mathsf{P}[t] \rightarrowtail s_t$ for every $t \in \mathsf{Tid}$; and $G_t \in G^t(s_t)$ for every $t \in \mathsf{Tid}$.

We note $C \rightarrowtail s$ to relate sequential program C with a possible sequence s of labels that it generates. This construction, defined in Fig. 23, is standard. Most operations are transformed as expected. Note that a fail compare-and-set produces both a fence and a read label. The new remote operations put/get generate two labels, since they perform a read followed by a write. For instance, we have $x := y^{\overline{n}} \rightarrowtail \mathtt{nrR}(y^{\overline{n}}, v), \mathtt{nlW}(x, v, \overline{n})$ for any arbitrary value v.

C.3 Extension of TSO Declarative Semantics

Definition C.1 (TSO-consistency). An execution $\langle E, \mathsf{po}, \mathsf{rf}\rangle$ is *TSO-consistent* (From [Lahav et al. 2016]) iff there exists a total store order tso such that:

(1) tso is a strict partial order on E that is total on $E.\mathcal{W}$;
(2) $\mathsf{ppo} \triangleq (\mathsf{po} \setminus E.\mathtt{lW} \times E.\mathtt{lR}) \subseteq \mathsf{tso}$;
(3) $\mathsf{rf_e} \triangleq (\mathsf{rf} \setminus \mathsf{po}) \subseteq \mathsf{tso}$; and
(4) if $(w, r) \in \mathsf{rf}$, $(w', r) \in (\mathsf{tso} \cup \mathsf{po})$, $w' \in \mathcal{W}$, and $\mathtt{loc}(w) = \mathtt{loc}(w')$, then $(w, w') \notin \mathsf{tso}$.

For programs without remote operations, the declarative semantics presented in Section 4 can be greatly simplified. The only relevant events are lW, lR, CAS, and F. In Fig. 9, only the first four lines and columns are useful, and we see that ippo simplifies to po, while oppo simplifies to $\mathsf{ppo} \triangleq \mathsf{po} \setminus (\mathtt{lW} \times \mathtt{lR})$. The relations pf and nfo are empty and no longer relevant.

Theorem C.2. *For programs without remote operations, RDMA^{TSO}-consistency (Def. 4.4) implies TSO-consistency (Def. C.1).*

Proof. After the simplifications detailed above, we assume $\mathsf{ib} = (\mathsf{po} \cup \mathsf{rf} \cup \mathsf{rb_b})^+$ and $\mathsf{ob} = (\mathsf{ppo} \cup \mathsf{rf_{\overline{b}}} \cup \mathsf{rb} \cup \mathsf{mo})^+$ to be irreflexive. ob is a strict partial order (irreflexive and transitive), and can be extended into a total order. So for tso we can choose any extension of ob that is total on $(\mathtt{lW} \cup \mathtt{CAS})$, satisfying the first point.

The second point comes directly from $\mathsf{ppo} \subseteq \mathsf{ob} \subseteq \mathsf{tso}$.

For the third point, it is enough to show $(\mathsf{rf} \setminus \mathsf{po}) \subseteq \mathsf{rf}_{\bar{\mathsf{b}}}$, since $\mathsf{rf}_{\bar{\mathsf{b}}} \subseteq \mathsf{ob} \subseteq \mathsf{tso}$. Since $\mathsf{rf}_{\bar{\mathsf{b}}} = \mathsf{rf} \setminus \mathsf{rf}_{\mathsf{b}}$, we want to show that $\mathsf{rf}_{\mathsf{b}} = [\mathtt{lW}];(\mathsf{rf} \cap \mathsf{sthd});[\mathtt{lR}] \subseteq \mathsf{po}$. Since $\mathsf{sthd} = \mathsf{po} \cup \mathsf{po}^{-1}$, it is enough to show that $\mathsf{rf} \cap \mathsf{po}^{-1}$ is empty, which comes from the fact that ib is irreflexive.

For the last point, let assume assume $(w, r) \in \mathsf{rf}$, $(w', r) \in (\mathsf{tso} \cup \mathsf{po})$, $w' \in \mathcal{W} = (\mathtt{lW} \cup \mathtt{CAS})$, and $\mathtt{loc}(w) = \mathtt{loc}(w')$. Since mo is total on write events on $\mathtt{loc}(w)$, either $(w, w') \in \mathsf{tso}$ or $(w', w) \in \mathsf{tso}$, and we want to show $(w', w) \in \mathsf{tso}$. Let us assume $(w, w') \in \mathsf{tso}$ and find a contradiction. By definition, we then have $r \xrightarrow{\mathsf{rb}} w'$.

If $(w', r) \in (\mathsf{po} \setminus \mathsf{ppo}) = [\mathtt{lW}]; \mathsf{po}; [\mathtt{lR}]$, then $r \xrightarrow{\mathsf{rb}_{\mathsf{b}}} w' \xrightarrow{\mathsf{po}} r$ which creates an illegal cycle in ib.

Otherwise, $(w', r) \in (\mathsf{tso} \cup \mathsf{ppo}) = \mathsf{tso}$, and we have $r \xrightarrow{\mathsf{rb}} w' \xrightarrow{\mathsf{tso}} r$ which creates an illegal cycle in tso since $\mathsf{rb} \subseteq \mathsf{ob} \subseteq \mathsf{tso}$. □

Note that the proof of theorem C.2 does not make use of the third condition of Def. 4.4. This is because it is redundant for programs without remote operations.

Theorem C.3. *For programs without remote operations, TSO-consistency (Def. C.1) implies* RDMA^{TSO}*-consistency (Def. 4.4).*

Proof. We assume $\langle E, \mathsf{po}, \mathsf{rf} \rangle$ and a relation tso satisfying the conditions of Def. C.1. We define $\mathsf{mo}_x \triangleq \mathsf{tso}|_{\mathcal{W}_x}$ total on writes on location x, and $\mathsf{mo} \triangleq \bigcup_{x \in \mathsf{Loc}} \mathsf{mo}_x$. Then $\langle E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \emptyset, \emptyset \rangle$ satisfies the conditions of a pre-execution (Def. 4.2). We now also have the derived relation rb, rf_{b}, $\mathsf{rf}_{\bar{\mathsf{b}}}$, and rb_{b}.

- We need to show that $\mathsf{ib} = (\mathsf{po} \cup \mathsf{rf} \cup \mathsf{rb}_{\mathsf{b}})^+$ is irreflexive.
  First, let us show that $\mathsf{rb}_{\mathsf{b}} = [\mathtt{lR}];(\mathsf{rb} \cap \mathsf{sthd});[\mathtt{lW}] \subseteq \mathsf{po}$. It is enough to check that $\mathsf{rb} \cap \mathsf{po}^{-1}$ is empty. By contradiction, if $r \xrightarrow{\mathsf{rb}} w' \xrightarrow{\mathsf{po}} r$, then by definition of rb there exists w such that $\mathtt{loc}(w) = \mathtt{loc}(w')$, $(w, r) \in \mathsf{rf}$, and $(w, w') \in \mathsf{mo} \subseteq \mathsf{tso}$, which contradicts TSO-consistency condition 4.
  It is then enough to show that $(\mathsf{po} \cup \mathsf{rf}_{\mathsf{e}})^+$ is irreflexive. By contradiction, let us assume a minimal cycle in $(\mathsf{po} \cup \mathsf{rf}_{\mathsf{e}})$. Since po is transitive, we can assume there is no two consecutive po edges in this cycle. Thus the left-side event of each po edge corresponds to the right-side event of an rf edge, and cannot be a local write event ($\mathtt{lW}$). Thus each po edge is also a ppo edge, and we have a cycle in $(\mathsf{ppo} \cup \mathsf{rf}_{\mathsf{e}}) \subseteq \mathsf{tso}$ (by TSO-consistency conditions 2 and 3), which is not allowed (by TSO-consistency condition 1).
- We need to show that $\mathsf{ob} = (\mathsf{ppo} \cup \mathsf{rf}_{\bar{\mathsf{b}}} \cup \mathsf{rb} \cup \mathsf{mo})^+$ is irreflexive.
  First, let us show $\mathsf{ob} \subseteq (\mathsf{tso} \cup \mathsf{rb})^+$. Since $\mathsf{ppo} \subseteq \mathsf{tso}$ (by TSO-consistency condition 2), $\mathsf{rf}_{\mathsf{e}} = (\mathsf{rf} \setminus \mathsf{po}) \subseteq \mathsf{tso}$ (by TSO-consistency condition 3), and $\mathsf{mo} \subseteq \mathsf{tso}$ (by definition of mo), it is enough to show that $\mathsf{rf}_{\bar{\mathsf{b}}} \subseteq (\mathsf{ppo} \cup \mathsf{rf}_{\mathsf{e}})$, i.e., that $(\mathsf{rf}_{\bar{\mathsf{b}}} \cap \mathsf{po}) \subseteq \mathsf{ppo}$. By definition, $\mathsf{rf}_{\bar{\mathsf{b}}} = \mathsf{rf} \setminus ([\mathtt{lW}];(\mathsf{po} \cup \mathsf{po}^{-1});[\mathtt{lR}])$, so $(\mathtt{lW} \times \mathtt{lR}) \cap (\mathsf{rf}_{\bar{\mathsf{b}}} \cap \mathsf{po}) = \emptyset$, and $(\mathsf{rf}_{\bar{\mathsf{b}}} \cap \mathsf{po}) \subseteq \mathsf{ppo}$.
  Second, we show that $(\mathsf{tso} \cup \mathsf{rb})^+$ is irreflexive. By contradiction, let us assume a minimal cycle in $(\mathsf{tso} \cup \mathsf{rb})$, and that edges labelled rb are in fact in $(\mathsf{rb} \setminus \mathsf{tso})$.
  - Since tso is transitive and irreflexive, this minimal cycle has at least one rb edge, and never two consecutive tso edges.
  - It cannot contain two consecutive rb edges: If $r_1 \xrightarrow{\mathsf{rb}} w_1 \xrightarrow{\mathsf{rb}} w_2 \rightarrow \cdots r_1$ then either $(w_1, w_2) \in \mathsf{tso}$ and the second edge can be labelled tso, or $(w_2, w_1) \in \mathsf{tso}$ and we have a shorter cycle $w_1 \xrightarrow{\mathsf{rb}} w_2 \xrightarrow{\mathsf{tso}} w_1$.

  - It cannot contain two non-consecutive rb edges: If $r_1 \xrightarrow{\text{rb}} w_1 \xrightarrow{\text{tso}} r_2 \xrightarrow{\text{rb}} w_2 \rightarrow \cdots r_1$ then either $(w_1, w_2) \in \text{tso}$ and we have a shorter cycle $r_1 \xrightarrow{\text{rb}} w_1 \xrightarrow{\text{tso}} w_2 \rightarrow \cdots r_1$, or $(w_2, w_1) \in \text{tso}$ and $(w_2, r_2) \in \text{tso}$ and we have a shorter cycle $r_2 \xrightarrow{\text{rb}} w_2 \xrightarrow{\text{tso}} r_2$.

  Thus we can assume a cycle of the form $r \xrightarrow{\text{rb}} w' \xrightarrow{\text{tso}} r$. By definition of rb there exists w such that $\texttt{loc}(w) = \texttt{loc}(w')$, $(w, r) \in \text{rf}$, and $(w, w') \in \text{mo} \subseteq \text{tso}$, which contradicts TSO-consistency condition 4.
  So $(\text{tso} \cup \text{rb})^+$ is irreflexive and ob is irreflexive.

- Lastly, we need to show that $([\texttt{Inst}]; \text{ib}; \text{ob})^+$ is irreflexive.
  As seen before, ib can be simplified to $(\text{po} \cup \text{rf}_\text{e})^+$. Since po is transitive, we have $[\texttt{Inst}]; \text{ib} = \left([\texttt{Inst}]; \text{rf}_\text{e}^+; (\text{po}; \text{rf}_\text{e}^+)^*; \text{po}^?\right) \cup \left([\texttt{Inst}]; \text{po}; (\text{rf}_\text{e}^+; \text{po})^*; \text{rf}_\text{e}^*\right)$. As previously, each po edge cannot start with a local write ($\texttt{lW}$) and is also a ppo edge. Thus $[\texttt{Inst}]; \text{ib} \subseteq (\text{ppo} \cup \text{rf}_\text{e})^+ \subseteq \text{tso}$.
  Since $\text{ob} \subseteq (\text{tso} \cup \text{rb})^+$, then we also have $([\texttt{Inst}]; \text{ib}; \text{ob})^+ \subseteq (\text{tso} \cup \text{rb})^+$ irreflexive.

□

C.4 Equivalent Declarative Semantics

Equivalently, the consistency conditions of Def. 4.4 can be restated with recursive definitions of ib and ob as follows. In this case, $(\text{ob}; [\texttt{Inst}])$ is included in ib while $([\texttt{Inst}]; \text{ib})$ is included in ob, and there is no need for the third condition.

*Definition C.4 (recursive-*RDMA$^{\text{TSO}}$*-consistency).* An execution $\langle E, \text{po}, \text{rf}, \text{mo}, \text{pf}, \text{nfo}\rangle$ is *RDMA$^{\text{TSO}}$-consistent* iff (1) ib is irreflexive; and (2) ob is irreflexive where:

$$\text{ib} \triangleq \left(\text{ippo} \cup \text{rf} \cup \text{pf} \cup \text{nfo} \cup \text{rb}_\text{b} \cup (\text{ob}; [\texttt{Inst}])\right)^+ \qquad \text{('issued-before')}$$

$$\text{ob} \triangleq \left(\text{oppo} \cup \text{rf}_{\bar{\text{b}}} \cup ([\texttt{nlW}]; \text{pf}) \cup \text{nfo} \cup \text{rb} \cup \text{mo} \cup ([\texttt{Inst}]; \text{ib})\right)^+ \qquad \text{('observed-before')}$$

Let us show that the two consistency definitions are equivalent. The definitions of ib and ob above can more formally be defined using limits as follows, where ib^0 and ob^0 correspond to the non-recursive definition of Def. 4.4.

$$\begin{aligned}
\text{ib}^0 &\triangleq (\text{ippo} \;\cup\; \text{rf} \;\cup\; \text{pf} \;\cup\; \text{rb}_\text{b} \;\cup\; \text{nfo})^+ \\
\text{ob}^0 &\triangleq \left(\text{oppo} \;\cup\; \text{rf}_{\bar{\text{b}}} \;\cup\; [\texttt{nlW}]; \text{pf} \;\cup\; \text{rb} \;\cup\; \text{nfo} \;\cup\; \text{mo}\right)^+ \\
\text{ib}^{n+1} &\triangleq (\text{ib}^n \cup \text{ob}^n; [\texttt{Inst}])^+ \\
\text{ob}^{n+1} &\triangleq (\text{ob}^n \cup [\texttt{Inst}]; \text{ib}^n)^+ \\
\text{ib} &\triangleq \lim_{n \to \infty} \text{ib}^n \\
\text{ob} &\triangleq \lim_{n \to \infty} \text{ob}^n
\end{aligned}$$

Theorem C.5. *recursive-*RDMA$^{\text{TSO}}$*-consistency (Def. C.4) implies* RDMA$^{\text{TSO}}$*-consistency (Def. 4.4).*

Proof. By contradiction, this amounts to showing that a cycle in either ib^0, ob^0, or $([\texttt{Inst}]; \text{ib}^0; \text{ob}^0)^+$ implies a cycle in ib or ob. This is trivial since $\text{ib}^0 \subseteq \text{ib}$, $\text{ob}^0 \subseteq \text{ob}$, and $([\texttt{Inst}]; \text{ib}^0; \text{ob}^0)^+ \subseteq \text{ob}$ □

Theorem C.6. RDMA$^{\text{TSO}}$*-consistency (Def. 4.4) implies recursive-*RDMA$^{\text{TSO}}$*-consistency (Def. C.4).*

Proof. Let us assume that there is no cycle in ib^0, ob^0, and $([\texttt{Inst}]; \text{ib}^0; \text{ob}^0)^+$.
Let us define $\text{A}^n \triangleq ([\texttt{Inst}]; \text{ib}^n; [\texttt{Inst}] \;\cup\; [\texttt{Inst}]; \text{ob}^n; [\texttt{Inst}] \;\cup\; [\texttt{Inst}]; \text{ib}^n; \text{ob}^n; [\texttt{Inst}])^+$.
Since ib^n and ob^n are transitive, it is clear that a cycle in A^n implies a cycle in either ib^n, ob^n, or $([\texttt{Inst}]; \text{ib}^n; \text{ob}^n)^+$. A corollary is that there is no cycle in A^0.

By contradiction, let us assume there is a cycle in either ib, ob, or $\lim_{n\to\infty} \mathsf{A}^n$. Let n be the highest integer such that ib^n, ob^n, and A^n have no cycle.

If there is a cycle in $\mathsf{ib}^{n+1} \triangleq (\mathsf{ib}^n \cup \mathsf{ob}^n; [\texttt{Inst}])^+$, and there is no cycle in ib^n, then the cycle is of the form $(\mathsf{ib}^n)^?; (\mathsf{ob}^n; [\texttt{Inst}]; \mathsf{ib}^n)^*; \mathsf{ob}^n; [\texttt{Inst}]; (\mathsf{ib}^n)^? \subseteq \mathsf{A}^n$, so we have a contradiction. Similarly a cycle in ob^{n+1} leads to the same contradiction.

Then let us assume a cycle in A^{n+1}. As explained previously, this implies a cycle in either ib^{n+1}, ob^{n+1}, or $([\texttt{Inst}]; \mathsf{ib}^{n+1}; \mathsf{ob}^{n+1})^+$. The first two cases have been explored and lead to a contradiction, so we can assume a cycle in $(([\texttt{Inst}]; \mathsf{ib}^{n+1}); (\mathsf{ob}^{n+1}; [\texttt{Inst}]))^+$.

Using the fact that ib^n and ob^n are transitive, we have:

$$\begin{aligned}
[\texttt{Inst}]; \mathsf{ib}^{n+1} &= [\texttt{Inst}]; (\mathsf{ib}^n \cup \mathsf{ob}^n; [\texttt{Inst}])^+ \\
&= ([\texttt{Inst}]; \mathsf{ib}^n) \;\cup\; \left([\texttt{Inst}]; (\mathsf{ib}^n)^?; (\mathsf{ob}^n; [\texttt{Inst}]; \mathsf{ib}^n)^*; \mathsf{ob}^n; [\texttt{Inst}]; (\mathsf{ib}^n)^?\right) \\
&\subseteq ([\texttt{Inst}]; \mathsf{ib}^n) \;\cup\; \mathsf{A}^n; ([\texttt{Inst}]; \mathsf{ib}^n)^?
\end{aligned}$$

Similarly, $(\mathsf{ob}^{n+1}; [\texttt{Inst}]) \subseteq (\mathsf{ob}^n; [\texttt{Inst}]) \;\cup\; (\mathsf{ob}^n; [\texttt{Inst}])^?; \mathsf{A}^n$.

Plugging the two halves together and simplifying the different cases, we see that we have $([\texttt{Inst}]; \mathsf{ib}^{n+1}); (\mathsf{ob}^{n+1}; [\texttt{Inst}]) \subseteq \mathsf{A}^n$. Thus we have a cycle in A^n, which contradicts our hypothesis. □

Note that the proof does not rely on the content of ib^0 and ob^0, so this recursive equivalent definition is also available when we do not assume the PCIe guarantees.

Similarly, the following definition is also equivalent:

Definition C.7 (RDMA^{TSO}-consistency). An execution $\langle E, \mathsf{po}, \mathsf{rf}, \mathsf{mo}, \mathsf{pf}, \mathsf{nfo}\rangle$ is *RDMA^{TSO}-consistent* iff (1) ib is irreflexive; and (2) ob is irreflexive where:

$$\mathsf{ib} \triangleq \left(\mathsf{ippo} \cup \mathsf{rf} \cup \mathsf{pf} \cup \mathsf{nfo} \cup \mathsf{rb}_{\mathsf{b}}\right)^+ \qquad \text{(‘issued-before’)}$$

$$\mathsf{ob} \triangleq \left(\mathsf{oppo} \cup \mathsf{rf}_{\overline{\mathsf{b}}} \cup ([\texttt{nlW}]; \mathsf{pf}) \cup \mathsf{nfo} \cup \mathsf{rb} \cup \mathsf{mo} \cup ([\texttt{Inst}]; \mathsf{ib})\right)^+ \qquad \text{(‘observed-before’)}$$

I.e., we extend ob with $([\texttt{Inst}]; \mathsf{ib})$ but we do not extend ib. Clearly, Def. C.4 implies Def. C.7 which implies Def. 4.4, so the three definitions are equivalent.

In the rest of the appendices, we use the recursive definition of consistency (Def. C.4 above), as it is easier to manipulate stronger ib/ob relations with fewer conditions.

C.5 Counter-example without explicit NIC buffer order

As explained, an NIC local (resp. remote) read flushes previous NIC local (resp. remote) writes on the same buffer. In Section 4, we ask for an explicit ordering nfo. If we do note require an ordering, but simply assert that no interleaving is allowed, then we need rules saying that ib and ob have to agree in these cases:

$$\mathsf{ib} \triangleq \begin{pmatrix} \mathsf{ippo} \;\cup\; \mathsf{rf} \;\cup\; \mathsf{pf} \;\cup\; \mathsf{rb}_{\mathsf{b}} \;\cup\; \mathsf{ob}; [\texttt{Inst}] \\ \cup\; [\texttt{nlR}]; (\mathsf{ob} \cap \mathsf{sqp}); [\texttt{nlW}] \\ \cup\; [\texttt{nrR}]; (\mathsf{ob} \cap \mathsf{sqp}); [\texttt{nrW}] \end{pmatrix}^+ \qquad \textit{(issued before)}$$

$$\mathsf{ob} \triangleq \begin{pmatrix} \mathsf{oppo} \;\cup\; \mathsf{rf}_{\overline{\mathsf{b}}} \;\cup\; [\texttt{nlW}]; \mathsf{pf} \;\cup\; \mathsf{rb} \;\cup\; \mathsf{mo} \;\cup\; [\texttt{Inst}]; \mathsf{ib} \\ \cup\; [\texttt{nlW}]; (\mathsf{ib} \cap \mathsf{sqp}); [\texttt{nlR}] \\ \cup\; [\texttt{nrW}]; (\mathsf{ib} \cap \mathsf{sqp}); [\texttt{nrR}] \end{pmatrix}^+ \qquad \textit{(observed before)}$$

The problem is that the two consistency conditions (ib and ob irreflexive) do not seem to be enough to recover a valid operational semantics execution.

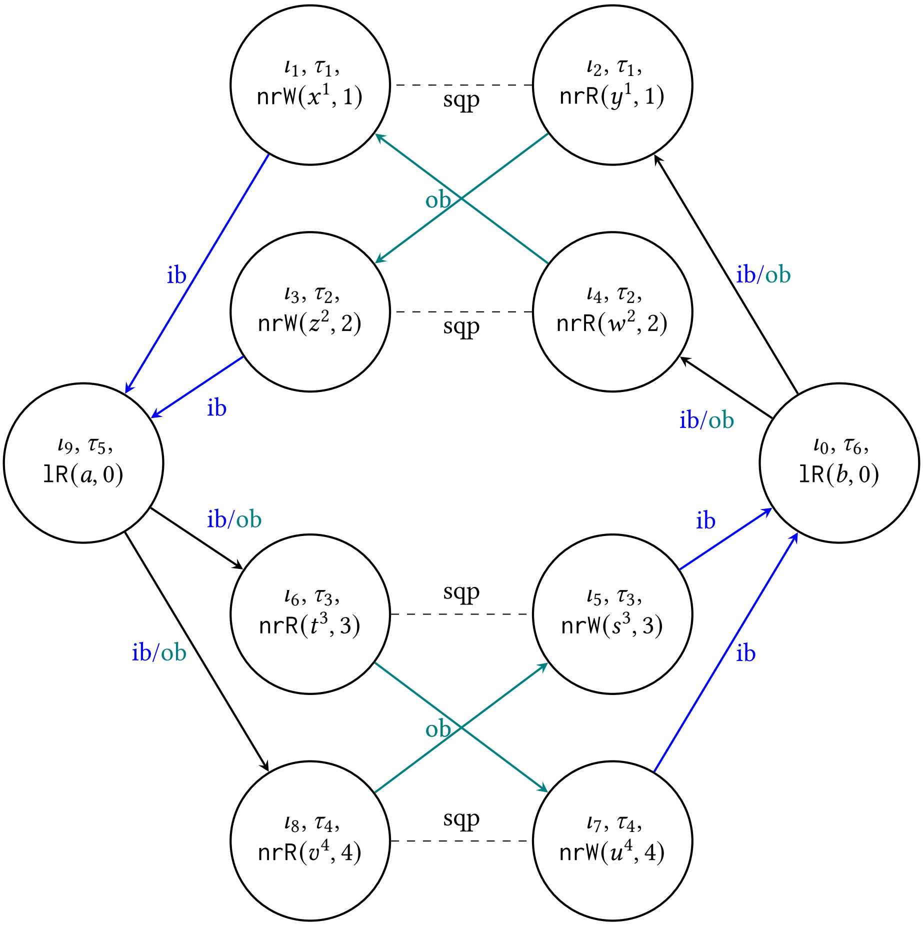
ι_1, τ_1, $\mathtt{nrW}(x^1, 1)$
sqp
ι_2, τ_1, $\mathtt{nrR}(y^1, 1)$
ob
ib
ib/ob
ι_3, τ_2, $\mathtt{nrW}(z^2, 2)$
sqp
ι_4, τ_2, $\mathtt{nrR}(w^2, 2)$
ib
ib/ob
ι_9, τ_5, $\mathtt{lR}(a, 0)$
ι_0, τ_6, $\mathtt{lR}(b, 0)$
ib/ob
ib
ι_6, τ_3, $\mathtt{nrR}(t^3, 3)$
sqp
ι_5, τ_3, $\mathtt{nrW}(s^3, 3)$
ib/ob
ib
ob
ι_8, τ_4, $\mathtt{nrR}(v^4, 4)$
sqp
ι_7, τ_4, $\mathtt{nrW}(u^4, 4)$

Fig. 24. Example of a consistent graph where ib and ob cannot be extended into total orders.

For instance, with the dummy example of Fig. 24, the consistency conditions are respected[5], but we quickly see that we cannot extend ib and ob into total orders respecting the recursive inclusions above.

Without loss of generality, if we decide that the first issued event is the one with label ι_1, i.e. we extend ib such that $\iota_1 \xrightarrow{\text{ib}} \iota_2$, then the same-queue-pair condition implies $\iota_1 \xrightarrow{\text{ob}} \iota_2$. By transitivity we have $\iota_4 \xrightarrow{\text{ob}} \iota_3$ which implies $\iota_4 \xrightarrow{\text{ib}} \iota_3$. On the other side of the graph, by transitivity we have $\iota_5 \xrightarrow{\text{ib}} \iota_6$ and $\iota_7 \xrightarrow{\text{ib}} \iota_8$, and the same-queue-pair condition forces $\iota_5 \xrightarrow{\text{ob}} \iota_6$ and $\iota_7 \xrightarrow{\text{ob}} \iota_8$. We thus have an ob cycle $\iota_5 \xrightarrow{\text{ob}} \iota_6 \xrightarrow{\text{ob}} \iota_7 \xrightarrow{\text{ob}} \iota_8 \xrightarrow{\text{ob}} \iota_5$.

While it is unclear if such a declarative semantics would allow for more behaviours than the operational semantics, we rely on the explicit ordering nfo to prove the equivalence between our two semantics.

[5] we should also add an arbitrary ordering on NIC remote writes because of mo, but it does not break consistency.

D ANNOTATED SEMANTICS

D.1 Annotated Labels and Inference Rules

On top of the 10 labels presented in Section 4, we create four new labels: $\mathtt{Put}(\overline{y}, x)$, $\mathtt{Get}(x, \overline{y})$, $\mathtt{nlEX}(\overline{n})$, and $\mathtt{nrEX}(\overline{n})$. These labels can also be used to create events (when bundled with an event identifier and a thread identifier).

We note E^{ext} the extended set of all events, including the four new labels.

Recall that $\mathcal{R} = \mathtt{lR} \cup \mathtt{CAS} \cup \mathtt{nlR} \cup \mathtt{nrR} \subseteq E^{\text{ext}}$ and $\mathcal{W} = \mathtt{lW} \cup \mathtt{CAS} \cup \mathtt{nlW} \cup \mathtt{nrW} \subseteq E^{\text{ext}}$. We also note $\mathtt{nEX} = \mathtt{nlEX} \cup \mathtt{nrEX}$.

$$
\begin{array}{lll}
\mathrm{ALabel} \ni \lambda \triangleq & \mid \mathtt{lR}\langle r, w\rangle & \text{where } r \in \mathtt{lR}, w \in \mathcal{W}, \mathsf{eq}_{\mathtt{loc\&v}}(r, w) \\
& \mid \mathtt{lW}\langle w\rangle & \text{where } w \in \mathtt{lW} \\
& \mid \mathtt{CAS}\langle u, w\rangle & \text{where } u \in \mathtt{CAS}, w \in \mathcal{W}, \mathsf{eq}_{\mathtt{loc\&v}}(u, w) \\
& \mid \mathtt{F}\langle f\rangle & \text{where } f \in \mathtt{F} \\
& \mid \mathtt{Push}\langle a\rangle & \text{where } a \in (\mathtt{Put} \cup \mathtt{Get} \cup \mathtt{nF}) \\
& \mid \mathtt{NIC}\langle a\rangle & \text{where } a \in (\mathtt{Put} \cup \mathtt{Get} \cup \mathtt{nF}) \\
& \mid \mathtt{nlR}\langle r, w, a, w'\rangle & \text{where } r \in \mathtt{nlR}, w \in \mathcal{W}, a \in \mathtt{Put}, w' \in \mathtt{nrW}, \mathsf{eq}_{\mathtt{loc\&v}}(r, w), \\
& & \qquad \mathtt{loc}_r(a) = \mathtt{loc}(r), \mathtt{loc}_w(a) = \mathtt{loc}(w'), v_{\mathrm{r}}(r) = v_{\mathrm{w}}(w') \\
& \mid \mathtt{nrR}\langle r, w, a, w'\rangle & \text{where } r \in \mathtt{nrR}, w \in \mathcal{W}, a \in \mathtt{Get}, w' \in \mathtt{nlW}, \mathsf{eq}_{\mathtt{loc\&v}}(r, w), \\
& & \qquad \mathtt{loc}_r(a) = \mathtt{loc}(r), \mathtt{loc}_w(a) = \mathtt{loc}(w'), v_{\mathrm{r}}(r) = v_{\mathrm{w}}(w') \\
& \mid \mathtt{nlW}\langle w, e\rangle & \text{where } w \in \mathtt{nlW}, e \in \mathtt{nlEX}, \mathsf{sameqp}(w, e) \\
& \mid \mathtt{nrW}\langle w, e\rangle & \text{where } w \in \mathtt{nrW}, e \in \mathtt{nrEX}, \mathsf{sameqp}(w, e) \\
& \mid \mathsf{CN}\langle e\rangle & \text{where } e \in \mathtt{nrEX} \\
& \mid \mathsf{P}\langle p, e\rangle & \text{where } p \in \mathsf{P}, e \in \mathtt{nEX}, \mathsf{sameqp}(p, e) \\
& \mid \mathsf{nF}\langle f\rangle & \text{where } f \in \mathsf{nF} \\
& \mid \mathsf{B}\langle w\rangle & \text{where } w \in \mathcal{W} \\
& \mid \mathcal{E}\langle t\rangle & \text{where } t \in \mathrm{Tid}
\end{array}
$$

$$
\begin{array}{rcl}
\mathsf{eq}_{\mathtt{loc\&v}}(r, w) & \triangleq & \mathtt{loc}(r) = \mathtt{loc}(w) \wedge v_{\mathrm{r}}(r) = v_{\mathrm{w}}(w) \\
\mathsf{sameqp}(e, e') & \triangleq & t(e) = t(e') \wedge \overline{n}(e) = \overline{n}(e')
\end{array}
$$

Fig. 25. Annotated Labels

For annotated labels, we reuse most names from labels, but they are different entities. For instance we note $r \in \mathtt{lR}$ for an event with label $\mathtt{lR}$, while $\lambda = \mathtt{lR}\langle \ldots \rangle$ is an annotated label.

We use type(λ) to denote the type of the annotated label ($\mathtt{lR}$, $\mathtt{lW}$, $\mathtt{CAS}$, F, $\mathtt{Push}$, $\mathtt{NIC}$, $\mathtt{nlR}$, $\mathtt{nrR}$, $\mathtt{nlW}$, $\mathtt{nrW}$, CN, P, nF, B, $\mathcal{E}$). We use $r(\lambda), w(\lambda), u(\lambda), a(\lambda), f(\lambda), p(\lambda), e(\lambda), \ldots$ to access the elements of a $\lambda \in$ ALabel where applicable. Also, we note $t(\lambda)$ for the thread of the first argument of λ.

The annotated program transitions (Fig. 26) use an additional annotated label $\mathtt{CASF}\langle r, w\rangle$ with $r \in \mathtt{lR}$ and $w \in \mathcal{W}$ to represent a failed CAS operation. This case is then translated into two labels (a memory fence and a local read) when creating a path in §D.2. Also, note that the annotated domains (e.g. the store buffers and the queue pairs) contain events, not annotated labels.

Program transitions: $\text{Prog} \xrightarrow{\text{ALabel}\uplus\{\mathsf{CASF}\}} \text{Prog}$ **Command transitions:** $\text{Comm} \xrightarrow{\text{ALabel}\uplus\{\mathsf{CASF}\}} \text{Comm}$

$$\frac{\mathsf{C}_1 \xrightarrow{\lambda} \mathsf{C}_1'}{\mathsf{C}_1;\mathsf{C}_2 \xrightarrow{\lambda} \mathsf{C}_1';\mathsf{C}_2} \qquad \frac{}{\mathsf{skip};\mathsf{C} \xrightarrow{\mathcal{E}\langle t\rangle} \mathsf{C}} \qquad \frac{i \in \{1,2\}}{\mathsf{C}_1 + \mathsf{C}_2 \xrightarrow{\mathcal{E}\langle t\rangle} \mathsf{C}_i^n} \qquad \frac{}{\mathsf{C}^* \xrightarrow{\mathcal{E}\langle t\rangle} \mathtt{skip}}$$

$$\frac{}{\mathsf{C}^* \xrightarrow{\mathcal{E}\langle t\rangle} \mathsf{C};\mathsf{C}^*} \qquad \frac{\mathsf{C} \rightsquigarrow \mathsf{C}'}{\mathsf{C} \xrightarrow{\mathcal{E}\langle t\rangle} \mathsf{C}'} \qquad \frac{\text{elocs}(e) = \emptyset \qquad w = (\iota, t, \mathtt{lW}(x, [\![e]\!]))}{x := e \xrightarrow{\mathtt{lW}\langle w\rangle} \mathtt{skip}}$$

$$\frac{\text{elocs}(e_{\text{old}}) = \text{elocs}(e_{\text{new}}) = \emptyset \qquad v \neq [\![e_{\text{old}}]\!] \qquad r = (\iota, t, \mathtt{lR}(x, v))}{z := \mathsf{CAS}(x, e_{\text{old}}, e_{\text{new}}) \xrightarrow{\mathsf{CASF}\langle r,w\rangle} z := v}$$

$$\frac{\text{elocs}(e_{\text{old}}) = \text{elocs}(e_{\text{new}}) = \emptyset \qquad u = (\iota, t, \mathsf{CAS}(x, [\![e_{\text{old}}]\!], [\![e_{\text{new}}]\!]))}{z := \mathsf{CAS}(x, e_{\text{old}}, e_{\text{new}}) \xrightarrow{\mathsf{CAS}\langle u,w\rangle} z := [\![e_{\text{old}}]\!]} \qquad \frac{f = (\iota, t, \mathsf{F})}{\mathtt{mfence} \xrightarrow{\mathsf{F}\langle f\rangle} \mathtt{skip}}$$

$$\frac{a = (\iota, t, \mathtt{Get}(x, \overline{y}))}{x := \overline{y} \xrightarrow{\mathsf{Push}\langle a\rangle} \mathtt{skip}} \qquad \frac{a = (\iota, t, \mathtt{Put}(\overline{y}, x))}{\overline{y} := x \xrightarrow{\mathsf{Push}\langle a\rangle} \mathtt{skip}} \qquad \frac{a = (\iota, t, \mathsf{nF}(\overline{n}))}{\mathtt{rfence}\ \overline{n} \xrightarrow{\mathsf{Push}\langle a\rangle} \mathtt{skip}} \qquad \frac{p = (\iota, t, \mathsf{P}(n))}{\mathtt{poll}(\overline{n}) \xrightarrow{\mathsf{P}\langle p,e\rangle} \mathtt{skip}}$$

$$\frac{r = (\iota, t, \mathtt{lR}(x, v))}{\mathtt{assume}(x = v) \xrightarrow{\mathtt{lR}\langle r,w\rangle} \mathtt{skip}} \qquad \frac{v \neq v' \qquad r = (\iota, t, \mathtt{lR}(x, v'))}{\mathtt{assume}(x \neq v) \xrightarrow{\mathtt{lR}\langle r,w\rangle} \mathtt{skip}} \qquad \frac{\mathsf{P}(t(\lambda)) \xrightarrow{\lambda} \mathsf{C}}{\mathsf{P} \xrightarrow{\lambda} \mathsf{P}[t(\lambda) \mapsto \mathsf{C}]}$$

Fig. 26. RDMA$^{\text{TSO}}$ program and command transitions for the annotated semantics

initialisation. Given a program P, let

$$\begin{array}{ll}
\mathsf{M}_0 \in \text{AMem} & \text{s.t. } \forall x \in \text{Loc. } \mathsf{M}_0(x) = init_x \text{ with } l(init_x) \triangleq \mathtt{lW}(x, 0) \\
\mathsf{b}_0 \in \text{ASBuff} & \mathsf{b}_0 \triangleq \varepsilon \\
\mathsf{B}_0 \in \text{ASBMap} & \mathsf{B}_0 \triangleq \lambda t.\mathsf{b}_0 \\
\mathsf{qp}_0 \in \text{AQPair} & \mathsf{qp}_0 \triangleq \langle \varepsilon, \varepsilon, \varepsilon \rangle \\
\mathsf{QP}_0 \in \text{AQPMap} & \mathsf{QP}_0 \triangleq \lambda t.\lambda n.\mathsf{qp}_0
\end{array}$$

$$\mathsf{M} \in \mathsf{AMem} \triangleq \{m \in \mathsf{Loc} \to \mathcal{W} \mid \forall x \in \mathsf{Loc}.\mathtt{loc}(m[x]) = x\} \qquad \mathsf{B} \in \mathsf{ASBMap} \triangleq \mathsf{Tid} \to \mathsf{ASBuff}$$

$$\mathsf{QP} \in \mathsf{AQPMap} \triangleq \mathsf{Tid} \to (\mathsf{Node} \to \mathsf{AQPair}) \qquad \mathsf{b} \in \mathsf{ASBuff} \triangleq (\mathtt{lR} \cup \mathtt{Get} \cup \mathtt{Put} \cup \mathtt{nF})^*$$

$$\mathbf{sqp} \in \mathsf{AQPair} \triangleq \mathsf{APipe} \times \mathsf{AWBR} \times \mathsf{AWBL} \qquad \mathbf{pipe} \in \mathsf{APipe} \triangleq (\mathtt{Get} \cup \mathtt{Put} \cup \mathtt{nF} \cup \mathtt{nrW} \cup \mathtt{nrEX} \cup \mathtt{nlW})^*$$

$$\mathbf{wb_R} \in \mathsf{AWBR} \triangleq \mathtt{nrW}^* \qquad \mathbf{wb_L} \in \mathsf{AWBR} \triangleq (\mathtt{nlW} \cup \mathtt{nlEX} \cup \mathtt{nrEX})^*$$

$$\frac{\mathsf{B}' = \mathsf{B}[t(w) \mapsto w \cdot \mathsf{B}(t(w))]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\mathtt{lW}\langle w \rangle} \mathsf{M}, \mathsf{B}', \mathsf{QP}} \qquad \frac{(\mathsf{M} \blacktriangleleft \mathsf{B}(t(r)))(\mathtt{loc}(r)) = w \qquad v_{\mathrm{r}}(r) = v_{\mathrm{W}}(w)}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\mathtt{lR}\langle r, w \rangle} \mathsf{M}, \mathsf{B}, \mathsf{QP}}$$

$$\frac{\mathsf{B}(t(u)) = \varepsilon \qquad \mathsf{M}(\mathtt{loc}(u)) = w \qquad v_{\mathrm{r}}(u) = v_{\mathrm{W}}(w)}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\mathtt{CAS}\langle u, w \rangle} \mathsf{M}[x \mapsto u], \mathsf{B}, \mathsf{QP}} \qquad \frac{\mathsf{B}(t(f)) = \varepsilon}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\mathtt{F}\langle f \rangle} \mathsf{M}, \mathsf{B}, \mathsf{QP}}$$

$$\frac{\mathsf{B}' = \mathsf{B}[t(a) \mapsto a \cdot \mathsf{B}(t(a))]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\mathsf{Push}\langle a \rangle} \mathsf{M}, \mathsf{B}', \mathsf{QP}} \qquad \frac{\mathsf{B}(t(w)) = \mathsf{b} \cdot w \qquad w \in \mathtt{lW}}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\mathsf{B}\langle w \rangle} \mathsf{M}[x \mapsto w], \mathsf{B}[t(w) \mapsto \mathsf{b}], \mathsf{QP}}$$

$$\frac{\mathsf{B}(t(a)) = \mathsf{b} \cdot a \qquad a \notin \mathtt{lW} \qquad \mathsf{QP}(t(a))(n(a)) = \mathsf{qp} \qquad \mathsf{qp}' = \mathsf{qp}[\mathbf{pipe} \mapsto a \cdot \mathsf{qp}.\mathbf{pipe}]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\mathtt{NIC}\langle a \rangle} \mathsf{M}, \mathsf{B}[t(a) \mapsto \mathsf{b}], \mathsf{QP}[t(a) \mapsto \mathsf{QP}(t(a))[n(a) \mapsto \mathsf{qp}']]}$$

$$\frac{\mathsf{QP}(t(p))(n(p)) = \mathsf{qp} \qquad \mathsf{qp}.\mathbf{wb_L} = \alpha \cdot e \qquad e \in \mathtt{nEX} \qquad \mathsf{qp}' = \mathsf{qp}[\mathbf{wb_L} \mapsto \alpha]}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\mathtt{P}\langle p, e \rangle} \mathsf{M}, \mathsf{B}, \mathsf{QP}[t(p) \mapsto \mathsf{QP}(t(p))[n(p) \mapsto \mathsf{qp}']]}$$

$$\frac{\mathsf{M}, \mathsf{QP}(t(\lambda))(\overline{n}) \xrightarrow{\lambda}_{\mathsf{sqp}} \mathsf{M}', \mathsf{qp}}{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\lambda} \mathsf{M}', \mathsf{B}, \mathsf{QP}[t(\lambda) \mapsto \mathsf{QP}(t(\lambda))[\overline{n} \mapsto \mathsf{qp}]]}$$

$$\text{with} \quad (\mathsf{M} \blacktriangleleft \alpha)(x) = \begin{cases} \mathsf{M}[x] & \alpha = \varepsilon \\ w & \alpha = w \cdot \beta \land w \in \mathcal{W} \land \mathtt{loc}(w) = x \\ (\mathsf{M} \blacktriangleleft \beta)(x) & \alpha = e \cdot \beta \land (e \notin \mathcal{W} \lor \mathtt{loc}(e) \neq x) \end{cases}$$

Fig. 27. RDMA$^{\mathsf{TSO}}$ hardware domains and hardware transitions for the annotated semantics

$$\frac{\mathbf{pipe} = \alpha \cdot f \qquad f = (\iota, t, \mathsf{nF}(n))}{\mathsf{M}, \langle \mathbf{pipe}, \mathbf{wb_R}, \mathbf{wb_L} \rangle \xrightarrow{\mathsf{nF}\langle f \rangle}_{\mathsf{sqp}} \mathsf{M}, \langle \alpha, \mathbf{wb_R}, \mathbf{wb_L} \rangle}$$

$$\frac{\begin{array}{c}\mathbf{pipe} = \alpha \cdot a \cdot \beta \qquad a = (\iota_a, t, \mathtt{Put}(\overline{y}, x)) \qquad \mathsf{M}(x) = w \qquad r = (\iota_r, t, \mathtt{nlR}(x, v_{\mathrm{w}}(w), n(\overline{y}))) \\ w' = (\iota_{w'}, t, \mathtt{nrW}(\overline{y}, v_{\mathrm{w}}(w))) \qquad \beta \in (\mathtt{nrW} \cup \mathtt{Get} \cup \mathtt{nlW} \cup \mathtt{nrEX})^* \qquad \mathbf{wb_L} \in \mathsf{nEX}^*\end{array}}{\mathsf{M}, \langle \mathbf{pipe}, \mathbf{wb_R}, \mathbf{wb_L} \rangle \xrightarrow{\mathtt{nlR}\langle r, w, a, w' \rangle}_{\mathsf{sqp}} \mathsf{M}, \langle \alpha \cdot w' \cdot \beta, \mathbf{wb_R}, \mathbf{wb_L} \rangle}$$

$$\frac{\begin{array}{c}\mathbf{pipe} = \alpha \cdot w \cdot \beta \\ w = (\iota_w, t, \mathtt{nrW}(\overline{y}, v)) \qquad e = (\iota_e, t, \mathtt{nrEX}(n(\overline{y}))) \qquad \beta \in (\mathtt{Get} \cup \mathtt{nlW} \cup \mathtt{nrEX})^*\end{array}}{\mathsf{M}, \langle \mathbf{pipe}, \mathbf{wb_R}, \mathbf{wb_L} \rangle \xrightarrow{\mathtt{nrW}\langle w, e \rangle}_{\mathsf{sqp}} \mathsf{M}, \langle \alpha \cdot e \cdot \beta, w \cdot \mathbf{wb_R}, \mathbf{wb_L} \rangle}$$

$$\frac{\mathbf{wb_R} = \alpha \cdot w \qquad w \in \mathtt{nrW}}{\mathsf{M}, \langle \mathbf{pipe}, \mathbf{wb_R}, \mathbf{wb_L} \rangle \xrightarrow{\mathsf{B}\langle w \rangle}_{\mathsf{sqp}} \mathsf{M}[v_{\mathrm{w}}(w) \mapsto w], \langle \mathbf{pipe}, \alpha, \mathbf{wb_L} \rangle}$$

$$\frac{\begin{array}{c}\mathbf{pipe} = \alpha \cdot a \cdot \beta \qquad a = (\iota_a, t, \mathtt{Get}(x, \overline{y})) \qquad \mathsf{M}(\overline{y}) = w \qquad r = (\iota_r, t, \mathtt{nrR}(\overline{y}, v_{\mathrm{w}}(w))) \\ w' = (\iota_{w'}, t, \mathtt{nlW}(x, v_{\mathrm{w}}(w), n(\overline{y}))) \qquad \beta \in (\mathtt{Get} \cup \mathtt{nlW} \cup \mathtt{nrEX})^* \qquad \mathbf{wb_R} = \varepsilon\end{array}}{\mathsf{M}, \langle \mathbf{pipe}, \mathbf{wb_R}, \mathbf{wb_L} \rangle \xrightarrow{\mathtt{nrR}\langle r, w, a, w' \rangle}_{\mathsf{sqp}} \mathsf{M}, \langle \alpha \cdot w' \cdot \beta, \mathbf{wb_R}, \mathbf{wb_L} \rangle}$$

$$\frac{\mathbf{pipe} = \alpha \cdot w \qquad w = (\iota_w, t, \mathtt{nlW}(x, v, n)) \qquad e = (\iota_e, t, \mathtt{nlEX}(n))}{\mathsf{M}, \langle \mathbf{pipe}, \mathbf{wb_R}, \mathbf{wb_L} \rangle \xrightarrow{\mathtt{nlW}\langle w, e \rangle}_{\mathsf{sqp}} \mathsf{M}, \langle \alpha, \mathbf{wb_R}, e \cdot w \cdot \mathbf{wb_L} \rangle}$$

$$\frac{\mathbf{pipe} = \alpha \cdot e \qquad e \in \mathtt{nrEX}}{\mathsf{M}, \langle \mathbf{pipe}, \mathbf{wb_R}, \mathbf{wb_L} \rangle \xrightarrow{\mathsf{CN}\langle e \rangle}_{\mathsf{sqp}} \mathsf{M}, \langle \alpha, \mathbf{wb_R}, e \cdot \mathbf{wb_L} \rangle}$$

$$\frac{\mathbf{wb_L} = \alpha \cdot w \cdot \beta \qquad w \in \mathtt{nlW} \qquad \beta \in \mathsf{nEX}^*}{\mathsf{M}, \langle \mathbf{pipe}, \mathbf{wb_R}, \mathbf{wb_L} \rangle \xrightarrow{\mathsf{B}\langle w \rangle}_{\mathsf{sqp}} \mathsf{M}[v_{\mathrm{w}}(w) \mapsto w], \langle \mathbf{pipe}, \mathbf{wb_R}, \alpha \cdot \beta \rangle}$$

Fig. 28. Annotated 3 Buffers NIC Semantics

D.2 Paths, Gluing, and Other Definitions

We define a path as: $\pi \in \mathsf{Path} \triangleq (\mathsf{ALabel} \setminus \mathcal{E}\langle t \rangle)^*$

We define Annotated Operational Semantics Gluing with the following rules.

$$\frac{\mathsf{P} \xrightarrow{\mathcal{E}\langle t \rangle} \mathsf{P}'}{\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi \Rightarrow \mathsf{P}', \mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi}$$

$$\frac{\mathsf{P} \xrightarrow{\lambda} \mathsf{P}' \qquad \mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\lambda} \mathsf{M}', \mathsf{B}', \mathsf{QP}' \qquad \lambda \in (\mathtt{lR} \cup \mathtt{lW} \cup \mathtt{CAS} \cup \mathsf{F} \cup \mathsf{Push} \cup \mathsf{P}) \qquad \mathsf{fresh}(\lambda, \pi)}{\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi \Rightarrow \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}', \lambda \cdot \pi}$$

$$\frac{\mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\lambda} \mathsf{M}', \mathsf{B}', \mathsf{QP}' \qquad \lambda \in (\mathtt{NIC} \cup \mathtt{nlR} \cup \mathtt{nrR} \cup \mathtt{nlW} \cup \mathtt{nrW} \cup \mathsf{CN} \cup \mathsf{nF} \cup \mathsf{B}) \qquad \mathsf{fresh}(\lambda, \pi)}{\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi \Rightarrow \mathsf{P}, \mathsf{M}', \mathsf{B}', \mathsf{QP}', \lambda \cdot \pi}$$

$$\frac{\begin{array}{c}\mathsf{P} \xrightarrow{\mathsf{CASF}\langle r, w \rangle} \mathsf{P}' \qquad \lambda_1 = \mathsf{F}\langle (\iota, t(r), \mathsf{F}) \rangle \\ \lambda_2 = \mathtt{lR}\langle r, w \rangle \qquad \mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\lambda_1} \mathsf{M}, \mathsf{B}, \mathsf{QP} \xrightarrow{\lambda_2} \mathsf{M}, \mathsf{B}, \mathsf{QP} \qquad \mathsf{fresh}(\lambda_1, \pi) \qquad \mathsf{fresh}(\lambda_2, \pi)\end{array}}{\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi \Rightarrow \mathsf{P}', \mathsf{M}, \mathsf{B}, \mathsf{QP}, \lambda_2 \cdot \lambda_1 \cdot \pi}$$

Two annotated labels are non-conflicting ($\lambda_1 \bowtie \lambda_2$) if they are of a different type or if their relevant arguments are disjoints. An annotated label is fresh if it does not conflict with any previous annotated label.

$$\mathbf{Relevant} : \mathsf{ALabel} \longrightarrow 2^{E^{\mathrm{ext}}}$$

$$\begin{aligned}
\mathbf{Relevant}(\mathtt{lR}\langle r, \_\rangle) &\triangleq \{r\} & \mathbf{Relevant}(\mathtt{nlW}\langle w, e\rangle) &\triangleq \{w, e\} \\
\mathbf{Relevant}(\mathtt{lW}\langle w\rangle) &\triangleq \{w\} & \mathbf{Relevant}(\mathtt{nrW}\langle w, e\rangle) &\triangleq \{w, e\} \\
\mathbf{Relevant}(\mathtt{CAS}\langle u, \_\rangle) &\triangleq \{u\} & \mathbf{Relevant}(\mathsf{CN}\langle e\rangle) &\triangleq \{e\} \\
\mathbf{Relevant}(\mathsf{F}\langle f\rangle) &\triangleq \{f\} & \mathbf{Relevant}(\mathsf{P}\langle p, e\rangle) &\triangleq \{p, e\} \\
\mathbf{Relevant}(\mathsf{Push}\langle a\rangle) &\triangleq \{a\} & \mathbf{Relevant}(\mathsf{nF}\langle f\rangle) &\triangleq \{f\} \\
\mathbf{Relevant}(\mathtt{NIC}\langle a\rangle) &\triangleq \{a\} & \mathbf{Relevant}(\mathsf{B}\langle w\rangle) &\triangleq \{w\} \\
\mathbf{Relevant}(\mathtt{nlR}\langle r, \_, a, w'\rangle) &\triangleq \{r, a, w'\} & \mathbf{Relevant}(\mathcal{E}\langle \_\rangle) &\triangleq \{\} \\
\mathbf{Relevant}(\mathtt{nrR}\langle r, \_, a, w'\rangle) &\triangleq \{r, a, w'\}
\end{aligned}$$

$$\begin{aligned}
\lambda_1 \bowtie \lambda_2 &\triangleq \mathsf{type}(\lambda_1) \neq \mathsf{type}(\lambda_2) \ \vee \mathbf{Relevant}(\lambda_1) \cap \mathbf{Relevant}(\lambda_2) = \emptyset \\
\mathsf{fresh}(\lambda, \pi) &\triangleq \forall \lambda' \in \pi,\ \lambda \bowtie \lambda' \\
\mathsf{nodup}(\pi) &\triangleq \forall \pi_2, \lambda, \pi_1.\ \pi = \pi_2 \cdot \lambda \cdot \pi_1 \implies \mathsf{fresh}(\lambda, \pi_1)
\end{aligned}$$

Relevant(λ) are the arguments that are important to consider to avoid duplicating events. The excluded events are the write operations we lookup when reading. For instance:

- Having both $\mathtt{lR}\langle r_1, w\rangle$ and $\mathtt{lR}\langle r_2, w\rangle$ during an execution is fine, since w can be looked up any number of time.
- Having both $\mathtt{nlR}\langle r_1, w_1, a, e_1\rangle$ and $\mathtt{nlR}\langle r_2, w_2, a, e_2\rangle$ during an execution is problematic, since it means the put operation a is being run twice.

$$
\begin{array}{l}
\mathrm{complete}(\pi) \triangleq \forall a, w', e, r, w, f. \\
\qquad\qquad \begin{array}{rl}
 & \mathtt{lW}\langle w\rangle \in \pi \implies \mathtt{B}\langle w\rangle \in \pi \\
\wedge & \mathtt{Push}\langle a\rangle \in \pi \implies \mathtt{NIC}\langle a\rangle \in \pi \\
\wedge & \mathtt{NIC}\langle f\rangle \in \pi \wedge f \in \mathtt{nF} \implies \mathtt{nF}\langle f\rangle \in \pi \\
\wedge & \mathtt{NIC}\langle a\rangle \in \pi \wedge a \in \mathtt{Put} \implies \exists r, w, w'.\ \mathtt{nlR}\langle r, w, a, w'\rangle \in \pi \\
\wedge & \mathtt{NIC}\langle a\rangle \in \pi \wedge a \in \mathtt{Get} \implies \exists r, w, w'.\ \mathtt{nrR}\langle r, w, a, w'\rangle \in \pi \\
\wedge & \mathtt{nlR}\langle r, w, a, w'\rangle \in \pi \implies \exists e.\ \mathtt{nrW}(w', e) \in \pi \\
\wedge & \mathtt{nrR}\langle r, w, a, w'\rangle \in \pi \implies \exists e.\ \mathtt{nlW}(w', e) \in \pi \\
\wedge & \mathtt{nlW}\langle w, e\rangle \in \pi \implies \mathtt{B}\langle w\rangle \in \pi \\
\wedge & \mathtt{nrW}\langle w, e\rangle \in \pi \implies \mathtt{B}\langle w\rangle \in \pi \wedge \mathtt{CN}\langle e\rangle \in \pi
\end{array}
\end{array}
$$

Informal: every pending operation is done and (most) buffers are empty. Note that some nEX (i.e., completion notifications) might still be in $\mathbf{wb}_{\mathrm{L}}$.

For a path π without duplicate (e.g. if nodup(π) holds), we define the total ordering of its annotated labels as follows. Note that the early part of the path is on the right.

$$\lambda_1 \prec_\pi \lambda_2 \ \triangleq\ \exists \pi_1, \pi_2, \pi_3 \text{ s.t. } \pi = \pi_3 \cdot \lambda_2 \cdot \pi_2 \cdot \lambda_1 \cdot \pi_1$$

backward completeness, with ordering.

$$
\begin{array}{l}
\mathrm{backComp}(\pi) \triangleq \forall a, w', e, r, w, f, p. \\
\qquad\qquad \begin{array}{rl}
 & \mathtt{B}\langle w\rangle \in \pi \implies \begin{pmatrix} \mathtt{lW}\langle w\rangle \prec_\pi \mathtt{B}\langle w\rangle \\ \vee\ \exists e. \mathtt{nlW}\langle w, e\rangle \prec_\pi \mathtt{B}\langle w\rangle \\ \vee\ \exists e. \mathtt{nrW}\langle w, e\rangle \prec_\pi \mathtt{B}\langle w\rangle \end{pmatrix} \\
\wedge & \mathtt{NIC}\langle a\rangle \in \pi \implies \mathtt{Push}\langle a\rangle \prec_\pi \mathtt{NIC}\langle a\rangle \\
\wedge & \mathtt{nF}\langle f\rangle \in \pi \implies \mathtt{NIC}\langle f\rangle \prec_\pi \mathtt{nF}\langle f\rangle \\
\wedge & \mathtt{nlR}\langle r, w, a, w'\rangle \in \pi \implies \mathtt{NIC}\langle a\rangle \prec_\pi \mathtt{nlR}\langle r, w, a, w'\rangle \\
\wedge & \mathtt{nrR}\langle r, w, a, w'\rangle \in \pi \implies \mathtt{NIC}\langle a\rangle \prec_\pi \mathtt{nrR}\langle r, w, a, w'\rangle \\
\wedge & \mathtt{nrW}\langle w', e\rangle \in \pi \implies \exists r, w, a.\ \mathtt{nlR}\langle r, w, a, w'\rangle \prec_\pi \mathtt{nrW}\langle w', e\rangle \\
\wedge & \mathtt{nlW}\langle w', e\rangle \in \pi \implies \exists r, w, a.\ \mathtt{nrR}\langle r, w, a, w'\rangle \prec_\pi \mathtt{nrW}\langle w', e\rangle \\
\wedge & \mathtt{CN}\langle e\rangle \in \pi \implies \exists w.\ \mathtt{nrW}\langle w, e\rangle \prec_\pi \mathtt{CN}\langle e\rangle \\
\wedge & \mathtt{P}\langle p, e\rangle \in \pi \implies \begin{pmatrix} \exists w.\ \mathtt{nlW}\langle w, e\rangle \prec_\pi \mathtt{B}\langle w\rangle \prec_\pi \mathtt{P}\langle p, e\rangle \\ \vee\ \mathtt{CN}\langle e\rangle \prec_\pi \mathtt{P}\langle p, e\rangle \end{pmatrix}
\end{array}
\end{array}
$$

Flush order

$$
\begin{array}{l}
\mathrm{bufFlushOrd}(\pi) \ \triangleq \\
\quad \forall w_1, w_2 \in \mathtt{lW}.\ \begin{pmatrix} t(w_1) = t(w_2) \implies \\ (\mathtt{B}\langle w_2\rangle \in \pi \wedge \mathtt{lW}\langle w_1\rangle \prec_\pi \mathtt{lW}\langle w_2\rangle) \iff \mathtt{B}\langle w_1\rangle \prec_\pi \mathtt{B}\langle w_2\rangle \end{pmatrix} \\
\wedge\ \forall a_1, a_2 \in (\mathtt{Get} \cup \mathtt{Put} \cup \mathtt{nF}).\ \begin{pmatrix} t(a_1) = t(a_2) \implies \\ (\mathtt{NIC}\langle a_2\rangle \in \pi \wedge \mathtt{Push}\langle a_1\rangle \prec_\pi \mathtt{Push}\langle a_2\rangle) \iff \mathtt{NIC}\langle a_1\rangle \prec_\pi \mathtt{NIC}\langle a_2\rangle \end{pmatrix}
\end{array}
$$

$$\land\ \forall a_1 \in (\mathtt{Get} \cup \mathtt{Put} \cup \mathtt{nF}), w_2 \in \mathtt{lW}.\ \left(\begin{array}{l} t(a_1) = t(w_2) \implies \\ \quad (\mathsf{B}\langle w_2\rangle \in \pi \land \mathsf{Push}\langle a_1\rangle \prec_\pi \mathtt{lW}\langle w_2\rangle) \iff \mathtt{NIC}\langle a_1\rangle \prec_\pi \mathsf{B}\langle w_2\rangle \\ \land\, (\mathtt{NIC}\langle a_1\rangle \in \pi \land \mathtt{lW}\langle w_2\rangle \prec_\pi \mathsf{Push}\langle a_1\rangle) \iff \mathsf{B}\langle w_2\rangle \prec_\pi \mathtt{NIC}\langle a_1\rangle \end{array}\right)$$

$$\land\ \forall w_1, w_2 \in \mathtt{nlW}.\ \left(\begin{array}{l} \mathsf{sameqp}(w_1, w_2) \implies \\ (\mathsf{B}\langle w_2\rangle \in \pi \land \mathtt{nlW}\langle w_1\rangle \prec_\pi \mathtt{nlW}\langle w_2\rangle) \iff \mathsf{B}\langle w_1\rangle \prec_\pi \mathsf{B}\langle w_2\rangle \end{array}\right)$$

$$\land\ \forall w_1, w_2 \in \mathtt{nrW}.\ \left(\begin{array}{l} \mathsf{sameqp}(w_1, w_2) \implies \\ (\mathsf{B}\langle w_2\rangle \in \pi \land \mathtt{nrW}\langle w_1\rangle \prec_\pi \mathtt{nrW}\langle w_2\rangle) \iff \mathsf{B}\langle w_1\rangle \prec_\pi \mathsf{B}\langle w_2\rangle \end{array}\right)$$

$$\land\ \forall w \in \mathtt{lW}, f \in \mathtt{F}.\ \mathtt{lW}\langle w\rangle \prec_\pi \mathtt{F}\langle f\rangle \land t(w) = t(f) \implies \mathsf{B}\langle w\rangle \prec_\pi \mathtt{F}\langle f\rangle$$

$$\land\ \forall w \in \mathtt{lW}, u \in \mathtt{CAS}.\ \mathtt{lW}\langle w\rangle \prec_\pi \mathtt{CAS}\langle u, \_\rangle \land t(w) = t(u) \implies \mathsf{B}\langle w\rangle \prec_\pi \mathtt{CAS}\langle u, \_\rangle$$

$$\land\ \forall w \in \mathtt{nlW}, r \in \mathtt{nlR}.\ (\mathtt{nlW}\langle w, \_\rangle \prec_\pi \mathtt{nlR}\langle r, \_, \_, \_\rangle \land \mathsf{sameqp}(w, r)) \implies \mathsf{B}\langle w\rangle \prec_\pi \mathtt{nlR}\langle r, \_, \_, \_\rangle$$

$$\land\ \forall w \in \mathtt{nrW}, r \in \mathtt{nrR}.\ (\mathtt{nrW}\langle w, \_\rangle \prec_\pi \mathtt{nrR}\langle r, \_, \_, \_\rangle \land \mathsf{sameqp}(w, r)) \implies \mathsf{B}\langle w\rangle \prec_\pi \mathtt{nrR}\langle r, \_, \_, \_\rangle$$

Poll order

$$\mathrm{pollOrder}(\pi) \triangleq \forall e_1, e_2.\ \left(\begin{array}{l} \quad \mathsf{sameqp}(e_1, e_2) \\ \land\, \lambda_1 \in \{\mathtt{nlW}\langle \_, e_1\rangle, \mathsf{CN}\langle e_1\rangle\} \\ \land\, \lambda_2 \in \{\mathtt{nlW}\langle \_, e_2\rangle, \mathsf{CN}\langle e_2\rangle\} \\ \land\, \lambda_1 \prec_\pi \lambda_2 \\ \land\, \mathsf{P}\langle \_, e_2\rangle \in \pi \end{array}\right) \implies \mathsf{P}\langle \_, e_1\rangle \prec_\pi \mathsf{P}\langle \_, e_2\rangle$$

NIC order

$$\begin{array}{l}
\mathrm{nicActOrder}(\pi) \triangleq \forall a_1, a_2.\ \mathtt{NIC}\langle a_1\rangle \prec_\pi \mathtt{NIC}\langle a_2\rangle \land \mathsf{sameqp}(a_1, a_2) \implies \\
\quad (a_1 \in \mathtt{nF} \land a_2 \in \mathtt{Get} \land \mathtt{nrR}\langle \_, \_, a_2, \_\rangle \in \pi \implies \mathtt{nF}\langle a_1\rangle \prec_\pi \mathtt{nrR}\langle \_, \_, a_2, \_\rangle) \\
\land\ (a_1 \in \mathtt{nF} \land a_2 \in \mathtt{Put} \land \mathtt{nlR}\langle \_, \_, a_2, \_\rangle \in \pi \implies \mathtt{nF}\langle a_1\rangle \prec_\pi \mathtt{nlR}\langle \_, \_, a_2, \_\rangle) \\
\land\ (a_1 \in \mathtt{nF} \land a_2 \in \mathtt{nF} \land \mathtt{nF}\langle a_2\rangle \in \pi \implies \mathtt{nF}\langle a_1\rangle \prec_\pi \mathtt{nF}\langle a_2\rangle) \\
\land\ (a_1 \in \mathtt{Get} \land a_2 \in \mathtt{nF} \land \mathtt{nF}\langle a_2\rangle \in \pi \implies \mathtt{nrR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \mathtt{nlW}\langle w_1, \_\rangle \prec_\pi \mathtt{nF}\langle a_2\rangle) \\
\land \left(\begin{array}{l} a_1 \in \mathtt{Put} \land a_2 \in \mathtt{nF} \land \mathtt{nF}\langle a_2\rangle \in \pi \\ \implies \mathtt{nlR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \mathtt{nrW}\langle w_1, e_1\rangle \prec_\pi \mathsf{CN}\langle e_1\rangle \prec_\pi \mathtt{nF}\langle a_2\rangle \end{array}\right) \\
\land \left(\begin{array}{l} a_1 \in \mathtt{Get} \land a_2 \in \mathtt{Get} \land \mathtt{nrR}\langle \_, \_, a_2, w_2\rangle \prec_\pi \mathtt{nlW}\langle w_2, \_\rangle \\ \implies \mathtt{nrR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \mathtt{nlW}\langle w_1, \_\rangle \prec_\pi \mathtt{nlW}\langle w_2, \_\rangle \end{array}\right) \\
\land \left(\begin{array}{l} a_1 \in \mathtt{Get} \land a_2 \in \mathtt{Put} \land \mathtt{nlR}\langle \_, \_, a_2, w_2\rangle \prec_\pi \mathtt{nrW}\langle w_2, e_2\rangle \prec_\pi \mathsf{CN}\langle e_2\rangle \\ \implies \mathtt{nrR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \mathtt{nlW}\langle w_1, \_\rangle \prec_\pi \mathsf{CN}\langle e_2\rangle \end{array}\right) \\
\land \left(\begin{array}{l} a_1 \in \mathtt{Put} \land a_2 \in \mathtt{Get} \land \mathtt{nrR}\langle \_, \_, a_2, \_\rangle \in \pi \\ \implies \mathtt{nlR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \mathtt{nrW}\langle w_1, \_\rangle \prec_\pi \mathtt{nrR}\langle \_, \_, a_2, \_\rangle \end{array}\right) \\
\land \left(\begin{array}{l} a_1 \in \mathtt{Put} \land a_2 \in \mathtt{Get} \land \mathtt{nrR}\langle \_, \_, a_2, w_2\rangle \prec_\pi \mathtt{nlW}\langle w_2, \_\rangle \\ \implies \mathtt{nlR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \mathtt{nrW}\langle w_1, e_1\rangle \prec_\pi \mathsf{CN}\langle e_1\rangle \prec_\pi \mathtt{nlW}\langle w_2\rangle \end{array}\right) \\
\land\ (a_1 \in \mathtt{Put} \land a_2 \in \mathtt{Put} \land \mathtt{nlR}\langle \_, \_, a_2, \_\rangle \in \pi \implies \mathtt{nlR}\langle \_, \_, a_1, \_\rangle \prec_\pi \mathtt{nlR}\langle \_, \_, a_2, \_\rangle) \\
\land \left(\begin{array}{l} a_1 \in \mathtt{Put} \land a_2 \in \mathtt{Put} \land \mathtt{nlR}\langle \_, \_, a_2, w_2\rangle \prec_\pi \mathtt{nrW}\langle w_2, \_\rangle \\ \implies \mathtt{nlR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \mathtt{nrW}\langle w_1, \_\rangle \prec_\pi \mathtt{nrW}\langle w_2, \_\rangle \end{array}\right) \\
\land \left(\begin{array}{l} a_1 \in \mathtt{Put} \land a_2 \in \mathtt{Put} \land \mathtt{nlR}\langle \_, \_, a_2, w_2\rangle \prec_\pi \mathtt{nrW}\langle w_2, e_2\rangle \prec_\pi \mathsf{CN}\langle e_2\rangle \\ \implies \mathtt{nlR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \mathtt{nrW}\langle w_1, e_1\rangle \prec_\pi \mathsf{CN}\langle e_1\rangle \prec_\pi \mathsf{CN}\langle e_2\rangle \end{array}\right)
\end{array}$$

Read order

$$
\begin{aligned}
\mathsf{wfrd}(\pi) \triangleq\quad & \forall \pi_2, r, w, \pi_1.\ \pi = \pi_2 \cdot \mathtt{lR}\langle r, w\rangle \cdot \pi_1 \implies \mathsf{wfrdCPU}(r, w, \pi_1) \\
\wedge\ & \forall \pi_2, u, w, \pi_1.\ \pi = \pi_2 \cdot \mathsf{CAS}\langle u, w\rangle \cdot \pi_1 \implies \mathsf{wfrdCPU}(u, w, \pi_1) \\
\wedge\ & \forall \pi_2, r, w, \pi_1.\ \pi = \pi_2 \cdot \mathtt{nlR}\langle r, w, \_, \_\rangle \cdot \pi_1 \implies \mathsf{wfrdNIC}(r, w, \pi_1) \\
\wedge\ & \forall \pi_2, r, w, \pi_1.\ \pi = \pi_2 \cdot \mathtt{nrR}\langle r, w, \_, \_\rangle \cdot \pi_1 \implies \mathsf{wfrdNIC}(r, w, \pi_1)
\end{aligned}
$$

$$
\begin{aligned}
\mathsf{wfrdCPU}(r, w, \pi) \triangleq\ & \left(\begin{array}{l}
\exists \pi_2, \lambda, \pi_1.\ \pi = \pi_2 \cdot \lambda \cdot \pi_1 \\
\wedge\ \lambda \in \{\mathsf{B}\langle w\rangle, \mathsf{CAS}\langle w, \_\rangle\} \\
\wedge\ \{\mathsf{B}\langle w'\rangle, \mathsf{CAS}\langle w', \_\rangle \in \pi_2 \mid \mathtt{loc}(w') = \mathtt{loc}(r)\} = \emptyset \\
\wedge\ \left\{ w' \;\middle|\; \begin{array}{l} \mathtt{lW}\langle w'\rangle \in \pi \wedge \mathsf{B}\langle w'\rangle \notin \pi\ \wedge \\ \mathtt{loc}(w') = \mathtt{loc}(r) \wedge t(w') = t(r) \end{array} \right\} = \emptyset
\end{array}\right) \\
\vee & \left(\begin{array}{l}
\exists \pi_2, \lambda, \pi_1.\ \pi = \pi_2 \cdot \lambda \cdot \pi_1 \\
\wedge\ \lambda = \mathtt{lW}\langle w\rangle \wedge t(w) = t(r) \wedge \mathsf{B}\langle w\rangle \notin \pi_2 \\
\wedge\ \{\mathtt{lW}\langle w'\rangle \in \pi_2 \mid \mathtt{loc}(w') = \mathtt{loc}(r) \wedge t(w') = t(r)\} = \emptyset
\end{array}\right) \\
\vee & \left(\begin{array}{l}
w = \mathit{init}_{\mathtt{loc}(w)} \wedge \\
\left\{ \begin{array}{l} \mathsf{B}\langle w'\rangle, \mathsf{CAS}\langle w', \_\rangle \in \pi, \\ \mathtt{lW}\langle w''\rangle \in \pi \end{array} \;\middle|\; \begin{array}{l} \mathtt{loc}(w') = \mathtt{loc}(r)\ \wedge \\ \mathtt{loc}(w'') = \mathtt{loc}(r) \wedge t(w'') = t(r) \end{array} \right\} = \emptyset
\end{array}\right)
\end{aligned}
$$

$$
\begin{aligned}
\mathsf{wfrdNIC}(r, w, \pi) \triangleq\ & \left(\begin{array}{l}
\exists \pi_2, \lambda, \pi_1.\ \pi = \pi_2 \cdot \lambda \cdot \pi_1 \\
\wedge\ \lambda \in \{\mathsf{B}\langle w\rangle, \mathsf{CAS}\langle w, \_\rangle\} \\
\wedge\ \{\mathsf{B}\langle w'\rangle, \mathsf{CAS}\langle w', \_\rangle \in \pi_2 \mid \mathtt{loc}(w') = \mathtt{loc}(r)\} = \emptyset
\end{array}\right) \\
\vee & \left(\begin{array}{l}
w = \mathit{init}_{\mathtt{loc}(w)} \wedge \\
\{\mathsf{B}\langle w'\rangle, \mathsf{CAS}\langle w', \_\rangle \in \pi \mid \mathtt{loc}(w') = \mathtt{loc}(r)\} = \emptyset
\end{array}\right)
\end{aligned}
$$

Well-formed path

$$
\begin{aligned}
\mathsf{wfp}(\pi) \triangleq\quad & \mathsf{nodup}(\pi) \\
\wedge\ & \mathsf{backComp}(\pi) \\
\wedge\ & \mathsf{bufFlushOrd}(\pi) \\
\wedge\ & \mathsf{pollOrder}(\pi) \\
\wedge\ & \mathsf{nicActOrder}(\pi) \\
\wedge\ & \mathsf{wfrd}(\pi)
\end{aligned}
$$

Definition D.1.

$$
\begin{aligned}
\mathsf{wf}(\mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi) \triangleq\quad & \mathsf{wfp}(\pi) \\
\wedge\ & \forall x \in \mathsf{Loc}.\ \mathsf{M}(x) = \mathsf{read}(\pi, x) \\
\wedge\ & \forall t \in \mathsf{Tid}.\mathsf{B}(t) = \mathsf{mksbuff}(\varepsilon, t, \pi) \\
\wedge\ & \forall t \in \mathsf{Tid}.\forall \overline{n} \in (\mathsf{Node} \setminus \{n(t)\}).\left(\begin{array}{l}
\mathsf{QP}(t)(\overline{n}).\mathbf{pipe} = \mathsf{mkpipe}(\varepsilon, t, \overline{n}, \pi) \\
\mathsf{QP}(t)(\overline{n}).\mathbf{wb_R} = \mathsf{mkwbR}(\varepsilon, t, \overline{n}, \pi) \\
\mathsf{QP}(t)(\overline{n}).\mathbf{wb_L} = \mathsf{mkwbL}(\varepsilon, t, \overline{n}, \pi)
\end{array}\right)
\end{aligned}
$$

Where, the functions read, mksbuff, mkpipe, mkwbR, and mkwbL are defined below.

$$\mathtt{read}(\lambda\cdot\pi, x) \triangleq \begin{cases} w & \lambda \in \{\mathtt{B}\langle w\rangle, \mathtt{CAS}\langle w, \_\rangle\} \land \mathtt{loc}(w) = x \\ \mathtt{read}(\pi, x) & \text{otherwise} \end{cases}$$

$$\mathtt{read}(\varepsilon, x) \triangleq \mathit{init}_x$$

$$\mathsf{mksbuff}(\mathsf{b}, t, \varepsilon) \triangleq \mathsf{b}$$

$$\mathsf{mksbuff}(\mathsf{b}, t, \pi\cdot\lambda) \triangleq \begin{cases} \mathsf{mksbuff}(w\cdot\mathsf{b}, t, \pi) & \lambda = \mathtt{lW}\langle w\rangle \land t(w) = t \land \mathtt{B}\langle w\rangle \notin \pi \\ \mathsf{mksbuff}(a\cdot\mathsf{b}, t, \pi) & \lambda = \mathtt{Push}\langle a\rangle \land \mathtt{NIC}\langle a\rangle \notin \pi \land t(a) = t \\ \mathsf{mksbuff}(\mathsf{b}, t, \pi) & \text{otherwise} \end{cases}$$

$$\mathsf{mkpipe}(\mathbf{pipe}, t, \overline{n}, \varepsilon) \triangleq \mathbf{pipe}$$

$$\mathsf{mkpipe}(\mathbf{pipe}, t, \overline{n}, \pi\cdot\lambda) \triangleq \begin{cases} \mathsf{mkpipe}(a\cdot\mathbf{pipe}, t, \overline{n}, \pi) & \text{if} \left(\begin{array}{l} t(\lambda) = t \land \ \overline{n}(\lambda) = \overline{n} \land \ \lambda = \mathtt{NIC}\langle a\rangle \\ \land\ \mathtt{nlR}\langle\_,\_, a, \_\rangle \notin \pi \land \ \mathtt{nrR}\langle\_,\_, a, \_\rangle \notin \pi \\ \land\ \mathtt{nF}\langle a\rangle \notin \pi \end{array} \right) \\ \mathsf{mkpipe}(w\cdot\mathbf{pipe}, t, \overline{n}, \pi) & \text{if} \left(\begin{array}{l} t(\lambda) = t \land \ \overline{n}(\lambda) = \overline{n} \land \ \lambda = \mathtt{NIC}\langle a\rangle \\ \land\ \mathtt{nlR}\langle\_,\_, a, w\rangle \in \pi \land \ \mathtt{nrW}\langle w, \_\rangle \notin \pi \end{array} \right) \\ \mathsf{mkpipe}(e\cdot\mathbf{pipe}, t, \overline{n}, \pi) & \text{if} \left(\begin{array}{l} t(\lambda) = t \land \ \overline{n}(\lambda) = \overline{n} \land \ \lambda = \mathtt{NIC}\langle a\rangle \\ \land\ \mathtt{nlR}\langle\_,\_, a, w\rangle \in \pi \land \ \mathtt{nrW}\langle w, e\rangle \in \pi \\ \land\ \mathtt{CN}\langle e\rangle \notin \pi \end{array} \right) \\ \mathsf{mkpipe}(w\cdot\mathbf{pipe}, t, \overline{n}, \pi) & \text{if} \left(\begin{array}{l} t(\lambda) = t \land \ \overline{n}(\lambda) = \overline{n} \land \ \lambda = \mathtt{NIC}\langle a\rangle \\ \land\ \mathtt{nrR}\langle\_,\_, a, w\rangle \in \pi \land \ \mathtt{nlW}\langle w, \_\rangle \notin \pi \end{array} \right) \\ \mathsf{mkpipe}(\mathbf{pipe}, t, \overline{n}, \pi) & \text{otherwise} \end{cases}$$

$$\mathsf{mkwbR}(\mathbf{wb_R}, t, \overline{n}, \varepsilon) \triangleq \mathbf{wb_R}$$

$$\mathsf{mkwbR}(\mathbf{wb_R}, t, \overline{n}, \pi\cdot\lambda) \triangleq \begin{cases} \mathsf{mkwbR}(w\cdot\mathbf{wb_R}, t, \overline{n}, \pi) & t(\lambda) = t \land \overline{n}(\lambda) = \overline{n} \land \lambda = \mathtt{nrW}\langle w, \_\rangle \land \mathtt{B}\langle w\rangle \notin \pi \\ \mathsf{mkwbR}(\mathbf{wb_R}, t, \overline{n}, \pi) & \text{otherwise} \end{cases}$$

$$\mathsf{mkwbL}(\mathbf{wb_L}, t, \overline{n}, \varepsilon) \triangleq \mathbf{wb_L}$$

$$\mathsf{mkwbL}(\mathbf{wb_L}, t, \overline{n}, \pi\cdot\lambda) \triangleq \begin{cases} \mathsf{mkwbL}(e\cdot w\cdot\mathbf{wb_L}, t, \overline{n}, \pi) & \text{if} \left(\begin{array}{l} t(\lambda) = t \land \ \overline{n}(\lambda) = \overline{n} \land \ \lambda = \mathtt{nlW}\langle w, e\rangle \\ \land\ \mathtt{B}\langle w\rangle \notin \pi \land \mathtt{P}\langle\_, e\rangle \notin \pi \end{array} \right) \\ \mathsf{mkwbL}(e\cdot\mathbf{wb_L}, t, \overline{n}, \pi) & \text{if} \left(\begin{array}{l} t(\lambda) = t \land \ \overline{n}(\lambda) = \overline{n} \land \ \lambda = \mathtt{nlW}\langle w, e\rangle \\ \land\ \mathtt{B}\langle w\rangle \in \pi \land \mathtt{P}\langle\_, e\rangle \notin \pi \end{array} \right) \\ \mathsf{mkwbL}(e\cdot\mathbf{wb_L}, t, \overline{n}, \pi) & \text{if} \left(\begin{array}{l} t(\lambda) = t \land \ \overline{n}(\lambda) = \overline{n} \land \ \lambda = \mathtt{CN}\langle e\rangle \\ \land\ \mathtt{P}\langle\_, e\rangle \notin \pi \end{array} \right) \\ \mathsf{mkwbL}(\mathbf{wb_L}, t, \overline{n}, \pi) & \text{otherwise} \end{cases}$$

Theorem D.2. *For all* $\mathsf{P}, \mathsf{P}', \mathsf{M}, \mathsf{M}', \mathsf{B}, \mathsf{B}', \mathsf{QP}, \mathsf{QP}', \pi, \pi'$*:*

- $\mathsf{wf}(\mathsf{M}_0, \mathsf{B}_0, \mathsf{QP}_0, \varepsilon)$*;*
- *if* $\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi \Rightarrow \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}', \pi'$ *and* $\mathsf{wf}(\mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi)$ *then* $\mathsf{wf}(\mathsf{M}', \mathsf{B}', \mathsf{QP}', \pi')$*;*

- *if* $\mathsf{P}, \mathsf{M}_0, \mathsf{B}_0, \mathsf{QP}_0, \varepsilon \Rightarrow^* (\lambda t.\mathtt{skip}), \mathsf{M}, \mathsf{B}_0, \mathsf{QP}, \pi$ *such that forall* $t, \overline{n}$ *we have* $\mathsf{QP}(t)(\overline{n}) = \langle \varepsilon, \varepsilon, \mathsf{nEX}^* \rangle$*, then* $\mathsf{wf}(\mathsf{M}, \mathsf{B}_0, \mathsf{QP}, \pi)$ *and* $\mathsf{complete}(\pi)$.

The proof of the first part follows trivially from the definitions of M_0, B_0, and QP_0. The second part is proved by induction on the structure of $\Rightarrow$. The last part follows from the previous two parts and induction on the length of $\Rightarrow^*$, as well as how the definition of wf on empty store buffers and queue pairs (regardless of nEX in $\mathbf{wb_L}$) implies $\mathsf{complete}(\pi)$.

D.3 From Annotated Semantics to Declarative Semantics

We define

$$\mathsf{getEG}(\pi) \triangleq \begin{cases} (\mathsf{Event}, \mathsf{po}, \mathsf{rf}, \mathsf{pf}, \mathsf{mo}, \mathsf{nfo}) & \text{if } \mathsf{wfp}(\pi) \land \mathsf{complete}(\pi) \\ \text{undefined} & \text{otherwise} \end{cases}$$

with

$$\mathsf{Event} \triangleq \mathsf{Event}_0 \cup \{\mathsf{getA}(\lambda) \mid \lambda \in \pi\}$$

Recall that Event_0 is the set of initialisation events $\{init_x \mid x \in \mathsf{Loc}\}$, where $l(init_x) = \mathtt{lW}(x, 0)$

$$\mathsf{getA} : \mathsf{ALabel} \rightharpoonup \mathsf{Event}$$

$$\begin{array}{rlrl}
\mathsf{getA}(\mathtt{lR}\langle r, \_\rangle) & \triangleq r & \mathsf{getA}(\mathtt{P}\langle p, \_\rangle) & \triangleq p \\
\mathsf{getA}(\mathtt{lW}\langle w\rangle) & \triangleq w & \mathsf{getA}(\mathtt{nF}\langle f\rangle) & \triangleq f \\
\mathsf{getA}(\mathtt{CAS}\langle u, \_\rangle) & \triangleq u & \mathsf{getA}(\mathtt{B}\langle w\rangle) & \triangleq w \\
\mathsf{getA}(\mathtt{F}\langle f\rangle) & \triangleq f & \mathsf{getA}(\mathtt{Push}\langle \_\rangle) & \text{is undefined} \\
\mathsf{getA}(\mathtt{nlR}\langle r, \_, \_, \_\rangle) & \triangleq r & \mathsf{getA}(\mathtt{NIC}\langle \_\rangle) & \text{is undefined} \\
\mathsf{getA}(\mathtt{nrR}\langle r, \_, \_, \_\rangle) & \triangleq r & \mathsf{getA}(\mathtt{CN}\langle \_\rangle) & \text{is undefined} \\
\mathsf{getA}(\mathtt{nlW}\langle w, \_\rangle) & \triangleq w & \mathsf{getA}(\mathcal{E}\langle \_\rangle) & \text{is undefined} \\
\mathsf{getA}(\mathtt{nrW}\langle w, \_\rangle) & \triangleq w & &
\end{array}$$

We define $\mathsf{getI}\lambda(\_, \pi)$ and $\mathsf{getO}\lambda(\_, \pi)$ to perform the reverse operation of getA. In the case of write events, $\mathsf{getI}\lambda(\_, \pi)$ retrieves the first label sending the write to the buffer, while $\mathsf{getO}\lambda(\_, \pi)$ retrieves the second label committing the write to memory.

$$\mathsf{getI}\lambda(\_, \pi), \mathsf{getO}\lambda(\_, \pi) : \{\mathsf{getA}(\lambda) \mid \lambda \in \pi\} \longrightarrow \mathsf{ALabel}$$

For all $\lambda \in \pi$:

- if $\mathsf{type}(\lambda) \in \{\mathtt{lR}, \mathtt{CAS}, \mathtt{F}, \mathtt{P}, \mathtt{nlR}, \mathtt{nrR}, \mathtt{nF}\}$, then $\mathsf{getI}\lambda(\mathsf{getA}(\lambda), \pi) \triangleq \mathsf{getO}\lambda(\mathsf{getA}(\lambda), \pi) \triangleq \lambda$;
- if $\mathsf{type}(\lambda) \in \{\mathtt{lW}, \mathtt{nlW}, \mathtt{nrW}\}$, then $\mathsf{getI}\lambda(\mathsf{getA}(\lambda), \pi) \triangleq \lambda$ while $\mathsf{getO}\lambda(\mathsf{getA}(\lambda), \pi) \triangleq \mathtt{B}\langle \mathtt{loc}(\lambda)\rangle$.
- if $\lambda = \mathtt{B}\langle w\rangle$, then from $\mathsf{backComp}(\pi)$ there is $\lambda' \prec_\pi \lambda$ such that $\mathsf{type}(\lambda) \in \{\mathtt{lW}, \mathtt{nlW}, \mathtt{nrW}\}$ and $\mathsf{getA}(\lambda') = \mathsf{getA}(\lambda) = w$. From the previous case, we have $\mathsf{getI}\lambda(w, \pi) \triangleq \lambda'$ and $\mathsf{getO}\lambda(w, \pi) \triangleq \lambda$.

From this we define two relations IB and OB on Event total on all meaningful events by copying the ordering in π.

$$\mathsf{IB} \triangleq \{(e_1, e_2) \mid \mathsf{getI}\lambda(e_1, \pi) \prec_\pi \mathsf{getI}\lambda(e_2, \pi)\} \cup (\mathsf{Event}_0 \times (\mathsf{Event} \setminus \mathsf{Event}_0))$$

$$\mathsf{OB} \triangleq \{(e_1, e_2) \mid \mathsf{getO}\lambda(e_1, \pi) \prec_\pi \mathsf{getO}\lambda(e_2, \pi)\} \cup (\mathsf{Event}_0 \times (\mathsf{Event} \setminus \mathsf{Event}_0))$$

From $\mathsf{wfp}(\pi)$, IB and OB are transitive and irreflexive. Note: we could make IB and OB total by adding an arbitrary total order on Event_0.

$$\mathsf{rf} \triangleq \{(w, r) \mid \mathtt{lR}\langle r, w\rangle \in \pi \lor \mathtt{nlR}\langle r, w, \_, \_\rangle \in \pi \lor \mathtt{nrR}\langle r, w, \_, \_\rangle \in \pi \lor \mathtt{CAS}\langle r, w\rangle \in \pi\}$$

$$\mathsf{pf} \triangleq \{(w, p) \mid \mathtt{nlW}\langle w, e\rangle \prec_\pi \mathtt{P}\langle p, e\rangle\} \cup \{(w, p) \mid \mathtt{nrW}\langle w, e\rangle \prec_\pi \mathtt{P}\langle p, e\rangle\}$$

$$\lambda \text{ generates } e \text{ in } \pi \;\triangleq\; \left(\begin{array}{c} \lambda \in \{\mathtt{lR}\langle e,\_\rangle, \mathtt{lW}\langle e\rangle, \mathtt{CAS}\langle e,\_\rangle, \mathtt{Push}\langle e\rangle, \mathtt{P}\langle e,\_\rangle, \mathtt{F}\langle e\rangle\} \\ \vee\, \lambda = \mathtt{Push}\langle a\rangle \wedge \left(\begin{array}{l} \lambda \prec_\pi \mathtt{nlR}\langle e,\_,a,\_\rangle \\ \vee\, \lambda \prec_\pi \mathtt{nlR}\langle \_,\_,a,e\rangle \\ \vee\, \lambda \prec_\pi \mathtt{nrR}\langle e,\_,a,\_\rangle \\ \vee\, \lambda \prec_\pi \mathtt{nrR}\langle \_,\_,a,e\rangle \end{array} \right) \end{array} \right)$$

$$\mathsf{po} \triangleq \mathsf{Event}_0 \times (\mathsf{Event} \setminus \mathsf{Event}_0) \cup \left\{ (e_1,e_2) \;\middle|\; \begin{array}{l} \lambda_1 \prec_\pi \lambda_2 \wedge t(\lambda_1) = t(\lambda_2) \\ \wedge\, \lambda_1 \text{ generates } e_1 \text{ in } \pi \\ \wedge\, \lambda_2 \text{ generates } e_2 \text{ in } \pi \end{array} \right\} \cup \left\{ (r,w) \;\middle|\; \begin{array}{l} \mathtt{nlR}\langle r,\_,\_,w\rangle \in \pi \\ \vee\, \mathtt{nrR}\langle r,\_,\_,w\rangle \in \pi \end{array} \right\}$$

$$\mathsf{mo} \triangleq \left\{ (w_1,w_2) \;\middle|\; \begin{array}{l} w_1 = init_x \\ \wedge\, (\mathtt{B}\langle w_2\rangle \in \pi \vee \mathtt{CAS}\langle w_2,\_\rangle \in \pi) \\ \wedge\, \mathtt{loc}(w_1) = x = \mathtt{loc}(w_2) \end{array} \right\} \cup \left\{ (w_1,w_2) \;\middle|\; \begin{array}{l} \lambda_1 \prec_\pi \lambda_2 \\ \wedge\, \lambda_1 \in \{\mathtt{B}\langle w_1\rangle, \mathtt{CAS}\langle w_1,\_\rangle\} \\ \wedge\, \lambda_2 \in \{\mathtt{B}\langle w_2\rangle, \mathtt{CAS}\langle w_2,\_\rangle\} \\ \wedge\, \mathtt{loc}(w_1) = \mathtt{loc}(w_2) \end{array} \right\}$$

$$\mathsf{nfo} \triangleq \left(\begin{array}{l} \{(r_1,w_2) \mid \mathsf{sameqp}(r_1,w_2) \wedge \mathtt{nlR}\langle r_1,\_,\_,\_\rangle \prec_\pi \mathtt{nlW}\langle w_2,\_\rangle \prec_\pi \mathtt{B}\langle w_2\rangle\} \\ \cup\; \{(r_1,w_2) \mid \mathsf{sameqp}(r_1,w_2) \wedge \mathtt{nrR}\langle r_1,\_,\_,\_\rangle \prec_\pi \mathtt{nrW}\langle w_2,\_\rangle \prec_\pi \mathtt{B}\langle w_2\rangle\} \\ \cup\; \{(w_1,r_2) \mid \mathsf{sameqp}(w_1,r_2) \wedge \mathtt{nlW}\langle w_1,\_\rangle \prec_\pi \mathtt{B}\langle w_1\rangle \prec_\pi \mathtt{nlR}\langle r_2,\_,\_,\_\rangle\} \\ \cup\; \{(w_1,r_2) \mid \mathsf{sameqp}(w_1,r_2) \wedge \mathtt{nrW}\langle w_1,\_\rangle \prec_\pi \mathtt{B}\langle w_1\rangle \prec_\pi \mathtt{nrR}\langle r_2,\_,\_,\_\rangle\} \end{array} \right)$$

From an execution graph $E = \mathrm{getEG}(\pi)$, we use the definitions of the paper to define oppo, ippo, rf_b, $\mathsf{rf}_{\overline{\mathsf{b}}}$, rb, rb_b, ob, and ib.

Theorem D.3. *$\mathrm{getEG}(\pi)$ is well-formed.*

Proof. We need to check the conditions of a pre-execution (Def. 4.2) and of well-formedness (Def. 4.3). For the pre-execution conditions:

- Checking $\mathsf{Event}^0 \times (\mathsf{Event} \setminus \mathsf{Event}^0) \subseteq \mathsf{po}$:
  by definition.
- Checking po is a union of strict partial orders each on one thread:
  If $t(\mathrm{e}_1) \neq t(\mathrm{e}_2)$, then $(\mathrm{e}_1,\mathrm{e}_2) \notin \mathsf{po}$ and $(\mathrm{e}_2,\mathrm{e}_1) \notin \mathsf{po}$ by definition. If $t(\mathrm{e}_1) = t(\mathrm{e}_2)$, then either $(\mathrm{e}_1,\mathrm{e}_2) \in \mathsf{po}$ or $(\mathrm{e}_2,\mathrm{e}_1) \in \mathsf{po}$. This comes from the second case of the definition of po: if there is λ_1 and λ_2 such that λ_i generates e_i in π, then either $\lambda_1 \prec_\pi \lambda_2$ or $\lambda_2 \prec_\pi \lambda_1$.
- Checking that rf is functional on its range:
  If $r \in \mathcal{R} \subseteq \{\mathrm{getA}(\lambda) \mid \lambda \in \pi\}$, then we have either $\mathtt{lR}\langle r,\_\rangle$, $\mathtt{nlR}\langle r,\_,\_,\_\rangle$, or $\mathtt{nrR}\langle r,\_,\_,\_\rangle$ in π, and r have at least one antecedent.
  If $(w,r) \in \mathsf{rf}$, let us assume $r \in \mathtt{nlR}$, then by definition $\mathtt{nlR}\langle r,w,\_,\_\rangle \in \pi$. Since $\mathrm{nodup}(\pi)$, for all $w' \neq w$, we have $\mathtt{nlR}\langle r,w',\_,\_\rangle \notin \pi$, and syntactically we cannot write $\mathtt{lR}\langle r,\_\rangle$ or $\mathtt{nrR}\langle r,\_,\_,\_\rangle$, so $(w',r) \notin \mathsf{rf}$. Similarly, $r \in \mathtt{lR}$ or $r \in \mathtt{nrR}$ only have one antecedent.
- Checking that rf relates events on the same location with matching values:
  By syntactic definition of the annotated labels lR, nlR, and nrR, e.g., $\mathtt{lR}\langle r,w\rangle \implies \mathsf{eq}_{\mathsf{loc\&v}}(r,w)$.
- Checking that mo is a union of strict total orders for writes on each variables:
  By definition of mo, given that we have $\mathrm{complete}(\pi)$, e.g., if $\mathtt{lW}\langle w\rangle \in \pi$ then $\mathtt{B}\langle w\rangle \in \pi$.
- Checking that $\mathsf{pf} \subseteq \mathsf{po} \cap \mathsf{sqp}$:
  If $(w,p) \in \mathsf{pf}$ with $w \in \mathtt{nlW}$ (resp. $\mathtt{nrW}$), then we have $\mathtt{nlW}\langle w,e\rangle \prec_\pi \mathtt{P}\langle p,e\rangle$. There is λ such that λ generates w in π, and we have $\lambda \prec_\pi \mathtt{nlW}\langle w,e\rangle \prec_\pi \mathtt{P}\langle p,e\rangle$. Also, $t(p) = t(w)$ and $\overline{n}(p) = \overline{n}(e) = \overline{n}(w)$, so we have $(w,p) \in \mathsf{po}$ and $(w,p) \in \mathsf{sqp}$.
- Checking that pf is functional on its domain:
  If $(w,p) \in \mathsf{pf}$ with $w \in \mathtt{nlW}$ (resp. $\mathtt{nrW}$), then we have $\mathtt{nlW}\langle w,e\rangle \prec_\pi \mathtt{P}\langle p,e\rangle$. From $\mathrm{nodup}(\pi)$, for all $p' \neq p$ we have $\mathtt{P}\langle p,e\rangle \notin \pi$, so w has at most one image.

- Checking that pf is total and functional on its range:
  If $p \in \mathsf{Event}$, then there is $e \in \texttt{nlEX}$ (resp. $\texttt{nrEX}$) such that $\texttt{P}\langle p, e\rangle \in \pi$. From $\mathrm{backComp}(\pi)$ there is $w \in \texttt{nlW}$ (resp. $\texttt{nrW}$) such that $\texttt{nlW}\langle w, e\rangle \prec_\pi \texttt{P}\langle p, e\rangle$, and so $(w, p) \in \mathsf{pf}$. From $\mathrm{nodup}(\pi)$, e cannot be used in another $\texttt{nlW}$ (resp. $\texttt{nrW}$) annotated label, and p has exactly one antecedent.
- Check that for all $(a, b) \in \mathsf{sqp}$, $a \in \texttt{nrR}$, $b \in \texttt{nrW}$, (resp. $\texttt{nlR}/\texttt{nlW}$) then $(a, b) \in \mathsf{nfo} \cup \mathsf{nfo}^{-1}$:
  By definition of nfo, given that $\mathrm{bufFlushOrd}(\pi)$ forbids the interleaving $\texttt{nrW}\langle w, \_\rangle \prec_\pi \texttt{nrR}\langle r, \_, \_, \_\rangle \prec_\pi \texttt{B}\langle w\rangle$ (resp. $\texttt{nlW}$ and $\texttt{nlR}$) when $\mathrm{sameqp}(r, w)$.

For the well-formedness conditions:

(1) Let us assume $(w_1, w_2) \in \mathsf{po} \cap \mathsf{sqp}$ and $(w_2, p_2) \in \mathsf{pf}$. The three events are on the same thread and queue pair.
If $w_1 \in \texttt{nlW}$, then by $\mathrm{complete}(\pi)$ there is a chain $\texttt{Push}\langle a_1\rangle \prec_\pi \texttt{NIC}\langle a_1\rangle \prec_\pi \texttt{nrR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \texttt{nlW}\langle w_1, e_1\rangle$; if $w_1 \in \texttt{nrW}$, there is instead a chain $\texttt{Push}\langle a_1\rangle \prec_\pi \texttt{NIC}\langle a_1\rangle \prec_\pi \texttt{nlR}\langle \_, \_, a_1, w_1\rangle \prec_\pi \texttt{nrW}\langle w_1, e_1\rangle \prec_\pi \texttt{CN}\langle e_1\rangle$. Similarly there is a chain for w_2. By $(w_1, w_2) \in \mathsf{po}$ we have $\texttt{Push}\langle a_1\rangle \prec_\pi \texttt{Push}\langle a_2\rangle$, and by $\mathrm{bufFlushOrd}(\pi)$ we have $\texttt{NIC}\langle a_1\rangle \prec_\pi \texttt{NIC}\langle a_2\rangle$.
Let us call λ_1 the last annotated label on the chain for w_1, i.e., either $\texttt{nlW}\langle w_1, e_1\rangle$ or $\texttt{CN}\langle e_1\rangle$. Similarly, λ_2 is the last annotated label on the chain for w_2. There is four cases to consider, but in all four $\mathrm{nicActOrder}(\pi)$ implies $\lambda_1 \prec_\pi \lambda_2$.
Then, from $\mathrm{pollOrder}(\pi)$, there is p_1 such that $\texttt{P}\langle p_1, e_1\rangle \prec_\pi \texttt{P}\langle p_2, e_2\rangle$. By definitions, we have both $(w_1, p_1) \in \mathsf{pf}$ and $(p_1, p_2) \in \mathsf{po}$.

(2) If $r \in \texttt{nlR}$, then there is $w \in \texttt{nrW}$ (taken from $\texttt{nlR}\langle r, \_, \_, w\rangle$) such that $(r, w) \in \mathsf{po}|_{\mathrm{imm}}$. This is by the last case of definition of po, since there is λ_a such that we have both λ_a generates r in π and λ_a generates w in π.
Similarly for $\texttt{nrR}/\texttt{nlW}$, $\texttt{nrW}/\texttt{nlR}$, and $\texttt{nlW}/\texttt{nrR}$.

(3) If $(r, w) \in \mathsf{po}|_{\mathrm{imm}}$, $\mathrm{type}(r) \in \{\texttt{nlR}, \texttt{nrR}\}$, and $\mathrm{type}(w) \in \{\texttt{nlW}, \texttt{nrW}\}$, then $(r, w) \in \mathsf{po}$ comes from the last case of the definition of po, and we have either $\texttt{nlR}\langle r, \_, \_, w\rangle$ or $\texttt{nrR}\langle r, \_, \_, w\rangle$ in π. In both cases, we have $v_{\mathrm{r}}(r) = v_{\mathrm{w}}(w)$ by syntactic definition of the annotated labels.

□

Lemma D.4. $\mathsf{OB}; [\texttt{Inst}] \subseteq \mathsf{IB}$ *and* $[\texttt{Inst}]; \mathsf{IB} \subseteq \mathsf{OB}$.

Proof. If $(e_1, e_2) \in \mathsf{OB}; [\texttt{Inst}]$, then $\mathrm{getO}\lambda(e_1, \pi) \prec_\pi \mathrm{getO}\lambda(e_2, \pi) = \mathrm{getI}\lambda(e_2, \pi)$.

- If $e_1 \in \texttt{Inst}$, then $\mathrm{getO}\lambda(e_1, \pi) = \mathrm{getI}\lambda(e_1, \pi)$, so we have $\mathrm{getI}\lambda(e_1, \pi) \prec_\pi \mathrm{getI}\lambda(e_2, \pi)$ and $(e_1, e_2) \in \mathsf{IB}$.
- If $e_1 \in \{\texttt{lW}, \texttt{nlW}, \texttt{nrW}\}$, there is λ such that $\mathrm{type}(\lambda) \in \{\texttt{lW}, \texttt{nlW}, \texttt{nrW}\}$, $\mathrm{getA}(\lambda) = e_1$, and $\mathrm{getI}\lambda(e_1, \pi) = \lambda \prec_\pi \texttt{B}\langle e_1\rangle = \mathrm{getO}\lambda(e_1, \pi)$. By transitivity we again have $\mathrm{getI}\lambda(e_1, \pi) \prec_\pi \mathrm{getI}\lambda(e_2, \pi)$ and $(e_1, e_2) \in \mathsf{IB}$.

With a similar reasoning, we can see that $[\texttt{Inst}]; \mathsf{IB} \subseteq \mathsf{OB}$.

□

Theorem D.5. $\mathrm{getEG}(\pi)$ *is consistent.*

Proof. From Definition C.4, we need to check that both ib and ob are irreflexive. Note that we still use the recursive definition given in Appendix C.4. Since IB and OB are irreflexive, it is enough to show that $\mathsf{ib} \subseteq \mathsf{IB}$ and $\mathsf{ob} \subseteq \mathsf{OB}$.

The explicit definition using limits is the following (where $\mathsf{rf}_{\overline{\mathsf{b}}} \triangleq (\mathsf{rf} \setminus \mathsf{rf}_{\mathsf{b}})$ includes $(\mathsf{rf} \cap \mathsf{sqp})$ since we assume the PCIe guarantees hold):

$$\mathsf{ib}^0 \triangleq (\mathsf{ippo} \ \cup \ \mathsf{rf} \ \cup \ \mathsf{pf} \ \cup \ \mathsf{rb}_{\mathsf{b}} \ \cup \ \mathsf{nfo})^+$$

$$
\begin{aligned}
\mathsf{ob}^0 &\triangleq (\mathsf{oppo} \,\cup\, \mathsf{rf}_{\bar{\mathsf{b}}} \,\cup\, [\mathtt{nlW}];\mathsf{pf} \,\cup\, \mathsf{rb} \,\cup\, \mathsf{nfo} \,\cup\, \mathsf{mo})^+ \\
\mathsf{ib}^{n+1} &\triangleq (\mathsf{ib}^n \cup \mathsf{ob}^n; [\mathtt{Inst}])^+ \\
\mathsf{ob}^{n+1} &\triangleq (\mathsf{ob}^n \cup [\mathtt{Inst}]; \mathsf{ib}^n)^+ \\
\mathsf{ib} &\triangleq \lim_{n\to\infty} \mathsf{ib}^n \\
\mathsf{ob} &\triangleq \lim_{n\to\infty} \mathsf{ob}^n
\end{aligned}
$$

It is then enough to show that $\mathsf{ib}^0 \subseteq \mathsf{IB}$ and $\mathsf{ob}^0 \subseteq \mathsf{OB}$. Using Lemma D.4 above, we can check the induction case:

$$
\begin{aligned}
\mathsf{ib}^{n+1} &= (\mathsf{ib}^n \cup \mathsf{ob}^n; [\mathtt{Inst}])^+ \subseteq (\mathsf{ib}^n \cup \mathsf{OB}; [\mathtt{Inst}])^+ \subseteq (\mathsf{IB} \cup \mathsf{IB})^+ = \mathsf{IB} \\
\mathsf{ob}^{n+1} &= (\mathsf{ob}^n \cup [\mathtt{Inst}]; \mathsf{ib}^n)^+ \subseteq (\mathsf{ob}^n \cup [\mathtt{Inst}]; \mathsf{IB})^+ \subseteq (\mathsf{OB} \cup \mathsf{OB})^+ = \mathsf{OB}
\end{aligned}
$$

Since IB and OB are transitive, we need to check the components of ib^0 and ob^0. There is eleven cases to verify.

- Checking $\mathsf{ippo} \subseteq \mathsf{IB}$.
  Let $E^{\mathrm{cpu}} = \{\mathtt{lR}, \mathtt{lW}, \mathtt{CAS}, \mathtt{F}, \mathtt{P}\}$ and $E^{\mathrm{nic}} = \{\mathtt{nlR}, \mathtt{nrR}, \mathtt{nlW}, \mathtt{nrW}, \mathtt{nF}\}$. $[E^{\mathrm{cpu}}]; \mathsf{po} \subseteq \mathsf{IB}$ by definition of po and IB: E^{cpu} are the events for which the same annotated label is used to define po and IB, i.e., $\forall e \in E^{\mathrm{cpu}}$, $\mathrm{getI}\lambda(e, \pi)$ generates e in π.
  To check that $[E^{\mathrm{nic}}]; \mathsf{ippo}; [E^{\mathrm{nic}}] \subseteq \mathsf{IB}$, there is 18 cases to consider. They are all trivially satisfied by $\mathrm{nicActOrder}(\pi)$.
- Checking $\mathsf{oppo} \subseteq \mathsf{OB}$.
  From above we have $[\mathtt{Inst}]; \mathsf{oppo} \subseteq [\mathtt{Inst}]; \mathsf{ippo} \subseteq [\mathtt{Inst}]; \mathsf{IB} \subseteq \mathsf{OB}$.
  $[\mathtt{lW}]; \mathsf{po}; [\mathrm{Event} \setminus (\mathtt{lR} \cup \mathtt{P})] \subseteq \mathsf{OB}$ by using $\mathrm{bufFlushOrd}(\pi)$.
  For the remaining four cases:

(I9) $[\mathtt{nlW}]; (\mathsf{po} \cap \mathsf{sqp}); [\mathtt{nlW}] \subseteq \mathsf{OB}$ comes from $\mathrm{nicActOrder}(\pi)$ (i.e., $\mathtt{nlW}\langle\ldots\rangle \prec_\pi \mathtt{nlW}\langle\ldots\rangle$) and $\mathrm{bufFlushOrd}(\pi)$ (i.e., $\mathtt{B}\langle\ldots\rangle \prec_\pi \mathtt{B}\langle\ldots\rangle$).

(G7) $[\mathtt{nrW}]; (\mathsf{po} \cap \mathsf{sqp}); [\mathtt{nrW}] \subseteq \mathsf{OB}$ comes from $\mathrm{nicActOrder}(\pi)$ (i.e., $\mathtt{nrW}\langle\ldots\rangle \prec_\pi \mathtt{nrW}\langle\ldots\rangle$) and $\mathrm{bufFlushOrd}(\pi)$ (i.e., $\mathtt{B}\langle\ldots\rangle \prec_\pi \mathtt{B}\langle\ldots\rangle$).

(G8) $[\mathtt{nrW}]; (\mathsf{po} \cap \mathsf{sqp}); [\mathtt{nrR}] \subseteq \mathsf{OB}$ comes from $\mathrm{nicActOrder}(\pi)$ (i.e., $\mathtt{nrW}\langle\ldots\rangle \prec_\pi \mathtt{nrR}\langle\ldots\rangle$) and $\mathrm{bufFlushOrd}(\pi)$ (i.e., $\mathtt{nrW}\langle\ldots\rangle \prec_\pi \mathtt{B}\langle\ldots\rangle \prec_\pi \mathtt{nrR}\langle\ldots\rangle$).

(G9) If $e_1 \in \mathtt{nrW}$, $e_3 \in \mathtt{nlW}$, and $(e_1, e_3) \in (\mathsf{po} \cap \mathsf{sqp})$, then from Def. 4.3 there is $e_2 \in \mathtt{nrR}$ such that $(e_2, e_3) \in \mathsf{po}|_{\mathrm{imm}}$ and thus $(e_1, e_2) \in (\mathsf{po} \cap \mathsf{sqp})$. From case G8 above, we have $(e_1, e_2) \in \mathsf{OB}$. From $\mathrm{backComp}(\pi)$, we have $(e_2, e_3) \in [\mathtt{Inst}]; \mathsf{IB} \subseteq \mathsf{OB}$. Thus $[\mathtt{nrW}]; (\mathsf{po} \cap \mathsf{sqp}); [\mathtt{nlW}] \subseteq \mathsf{OB}$.

- Checking $\mathsf{rf}_{\bar{\mathsf{b}}} \subseteq \mathsf{OB}$.
  If $(w, r) \in \mathsf{rf}_{\bar{\mathsf{b}}}$, there is π_1 and π_2 such that $\pi = \pi_2 \cdot \mathrm{getO}\lambda(r, \pi) \cdot \pi_1$, and we use $\mathrm{wfrd}(\pi)$.
  - If $r \in \mathtt{lR}$, we have $\mathrm{wfrdCPU}(r, w, \pi_1)$. The definition allow for three different cases. In the first case, $\lambda \in \{\mathtt{B}\langle w\rangle, \mathtt{CAS}\langle w, \_\rangle\}$ is in π_1; we have $\lambda = \mathrm{getO}\lambda(w, \pi) \prec_\pi \mathrm{getO}\lambda(r, \pi)$ and so $(w, r) \in \mathsf{OB}$. In the second case, we have $\lambda = \mathtt{lW}\langle w\rangle$ and $t(w) = t(r)$; so $(w, r) \in [\mathtt{lW}]; (\mathsf{rf} \cap \mathsf{sthd}); [\mathtt{lR}] = \mathsf{rf}_{\mathsf{b}}$, which contradicts $(w, r) \in \mathsf{rf}_{\bar{\mathsf{b}}} = \mathsf{rf} \setminus \mathsf{rf}_{\mathsf{b}}$. In the third case, $w = \mathit{init}_x$ for some location x, so $(w, r) \in \mathrm{Event}_0 \times (\mathrm{Event} \setminus \mathrm{Event}_0) \subseteq \mathsf{OB}$.
  - If $r \in \mathtt{CAS}$, similarly to above, except the second case of $\mathrm{wfrdCPU}(r, w, \pi_1)$ is not possible because of $\mathrm{bufFlushOrd}(\pi)$: $\mathtt{B}\langle w\rangle \notin \pi_1$ while $\mathtt{CAS}$ acts as a memory fence.
  - If $r \in \mathtt{nlR}$, we have $\mathrm{wfrdNIC}(r, w, \pi_1)$, with two possibilities. In the first case, $\lambda \in \{\mathtt{B}\langle w\rangle, \mathtt{CAS}\langle w, \_\rangle\}$ is in π_1; we have $\lambda = \mathrm{getO}\lambda(w, \pi) \prec_\pi \mathrm{getO}\lambda(r, \pi)$ and so $(w, r) \in \mathsf{OB}$. In the second case, $w = \mathit{init}_x$ for some location x, so $(w, r) \in \mathrm{Event}_0 \times (\mathrm{Event} \setminus \mathrm{Event}_0) \subseteq \mathsf{OB}$.
  - If $r \in \mathtt{nrR}$, similarly to above.

- Checking $\mathsf{rf} \subseteq \mathsf{IB}$.
  From above we have $\mathsf{rf}_{\overline{\mathsf{b}}} = \mathsf{rf}_{\overline{\mathsf{b}}}; [\mathtt{Inst}] \subseteq \mathsf{OB}; [\mathtt{Inst}] \subseteq \mathsf{IB}$.
  If $(w, r) \in \mathsf{rf}_{\mathsf{b}} \subseteq [\mathtt{lW}]; \mathsf{rf}; [\mathtt{lR}]$, then there is $\mathtt{lR}\langle r, w\rangle \in \pi$. There is π_1 and π_2 such that $\pi = \pi_2 \cdot \mathtt{lR}\langle r, w\rangle \cdot \pi_1$. So by $\mathrm{wfrd}(\pi)$ we have $\mathrm{wfrdCPU}(r, w, \pi_1)$ which implies $\mathtt{lW}\langle w\rangle \prec_\pi \mathtt{lR}\langle r, w\rangle$ and $(w, r) \in \mathsf{IB}$.
- Checking $[\mathtt{nlW}]; \mathsf{pf} \subseteq \mathsf{OB}$.
  If $(w, p) \in \mathsf{pf}$ with $w \in \mathtt{nlW}$, then there exists e such that $\mathtt{nlW}\langle w, e\rangle \prec_\pi \mathtt{P}\langle p, e\rangle$. From $\mathrm{backComp}(\pi)$, we have $\mathtt{nlW}\langle w, e\rangle \prec_\pi \mathtt{B}\langle w\rangle \prec_\pi \mathtt{P}\langle p, e\rangle$ and so $(w, p) \in \mathsf{OB}$.
- Checking $\mathsf{pf} \subseteq \mathsf{IB}$.
  If $(w, p) \in \mathsf{pf}$, then there exists e such that either $\mathtt{nlW}\langle w, e\rangle \prec_\pi \mathtt{P}\langle p, e\rangle$ or $\mathtt{nrW}\langle w, e\rangle \prec_\pi \mathtt{P}\langle p, e\rangle$. In both cases we immediately have $(w, p) \in \mathsf{IB}$.
- Checking $\mathsf{rb}_{\mathsf{b}} \subseteq \mathsf{IB}$.
  If $(r, w') \in \mathsf{rb}_{\mathsf{b}}$ then $r \in \mathtt{lR}$, $w' \in \mathtt{lW}$, $t(r) = t(w')$, and there exists w such that $(w, r) \in \mathsf{rf}$ and $(w, w') \in \mathsf{mo}$. There is π_4 and π_3 such that $\pi = \pi_4 \cdot \mathtt{lR}\langle r, w\rangle \cdot \pi_3$. So by $\mathrm{wfrd}(\pi)$ we have $\mathrm{wfrdCPU}(r, w, \pi_3)$, and there is three cases to consider.
  - In the first case, $\pi_3 = \pi_2 \cdot \lambda_w \cdot \pi_1$, with $\lambda_w \in \{\mathtt{B}\langle w\rangle, \mathtt{CAS}\langle w, \_\rangle\}$, and $\mathtt{B}\langle w'\rangle \notin \pi_2$. Since $(w, w') \in \mathsf{mo}$ we have $\mathtt{B}\langle w'\rangle \notin \pi_1$, an so $\mathtt{B}\langle w'\rangle \notin \pi_3$. The last condition of the first case then gives us $\mathtt{lW}\langle w'\rangle \notin \pi_3$, which implies $(r, w') \in \mathsf{IB}$.
  - In the second case, $\pi_3 = \pi_2 \cdot \lambda_w \cdot \pi_1$, with $\lambda_w = \mathtt{lW}\langle w\rangle$, $\mathrm{thread}(w) = \mathrm{thread}(r)$, and $\mathtt{B}\langle w\rangle \notin \pi_3$. Then w and w' are on the same thread, and by $\mathrm{bufFlushOrd}(\pi)$ and $(w, w') \in \mathsf{mo}$ we have $\mathtt{lW}\langle w\rangle \prec_\pi \mathtt{lW}\langle w'\rangle$ and $\mathtt{lW}\langle w'\rangle \notin \pi_1$. The last condition of the second case gives us $\mathtt{lW}\langle w'\rangle \notin \pi_2$, so $\mathtt{lW}\langle w'\rangle \notin \pi_3$ and $(r, w') \in \mathsf{IB}$.
  - In the last case, $w = \mathit{init}_x$ for some location x, and we immediately get $\mathtt{lW}\langle w'\rangle \notin \pi_3$, which implies $(r, w') \in \mathsf{IB}$.
- Checking $\mathsf{rb} \subseteq \mathsf{OB}$.
  If $(r, w') \in \mathsf{rb}$, then there exists w such that $(w, r) \in \mathsf{rf}$ and $(w, w') \in \mathsf{mo}$. By definition of rf, there is π_4 and π_3 such that $\pi = \pi_4 \cdot \lambda_r \cdot \pi_3$, with $\lambda_r \in \{\mathtt{lR}\langle r, w\rangle, \mathtt{CAS}\langle r, w\rangle, \mathtt{nlR}\langle r, w, \_, \_\rangle, \mathtt{nrR}\langle r, w, \_, \_\rangle\}$. So by $\mathrm{wfrd}(\pi)$ we have either $\mathrm{wfrdNIC}(r, w, \pi_3)$ or $\mathrm{wfrdCPU}(r, w, \pi_3)$, and there is five cases to consider.
  - In the first case of $\mathrm{wfrdNIC}(r, w, \pi_3)$, $\pi_3 = \pi_2 \cdot \mathrm{getO}\lambda(w, \pi) \cdot \pi_1$, and $\mathrm{getO}\lambda(w', \pi) \notin \pi_2$. Since $(w, w') \in \mathsf{mo}$ we have $\mathrm{getO}\lambda(w', \pi) \notin \pi_1$, and thus $\mathrm{getO}\lambda(w', \pi) \notin \pi_3$. So $\mathrm{getO}\lambda(w', \pi) \in \pi_4$ and $(r, w') \in \mathsf{OB}$.
  - In the last case $\mathrm{wfrdNIC}(r, w, \pi_3)$, $w = \mathit{init}_x$ for some location x, and we immediately have $\mathrm{getO}\lambda(w', \pi) \notin \pi_3$, which implies $(r, w') \in \mathsf{OB}$.
  - For the first case of $\mathrm{wfrdCPU}(r, w, \pi_3)$, same reasoning as for the first case of wfrdNIC.
  - For the second case of $\mathrm{wfrdCPU}(r, w, \pi_3)$, $\pi_3 = \pi_2 \cdot \mathrm{getI}\lambda(w, \pi) \cdot \pi_1$, with $\mathrm{thread}(w) = \mathrm{thread}(r)$, and $\mathrm{getO}\lambda(w, \pi) \notin \pi_3$. So $\mathrm{getO}\lambda(w, \pi) \in \pi_4$, and since $(w, w') \in \mathsf{mo}$ we have $\mathrm{getO}\lambda(w', \pi) \in \pi_4$ as well, and $(r, w') \in \mathsf{OB}$.
  - For the last case of $\mathrm{wfrdCPU}(r, w, \pi_3)$, same reasoning as for the last case of wfrdNIC.
- Checking $\mathsf{nfo} \subseteq \mathsf{IB}$.
  By definition of nfo.
- Checking $\mathsf{nfo} \subseteq \mathsf{OB}$.
  By definition of nfo.
- Checking $\mathsf{mo} \subseteq \mathsf{OB}$.
  By definition of mo, as what matters are the init_x, $\mathtt{B}\langle w\rangle$, and $\mathtt{CAS}\langle w, \_\rangle$ events.

□

D.4 From Declarative Semantics to Annotated Semantics

From a program P and a well-formed consistent execution graph $G = (\mathsf{Event}, \mathsf{po}, \mathsf{rf}, \mathsf{pf}, \mathsf{mo}, \mathsf{nfo})$, where $(\mathsf{Event}, \mathsf{po})$ is generated by P, we want to reconstruct an annotated semantics execution.

Theorem D.6. *ib and ob can be extended into total relations IB and OB on Event such that:*

- *IB and OB are irreflexive and transitive;*
- $\mathsf{OB}; [\mathtt{Inst}] \subseteq \mathsf{IB}$ *and* $[\mathtt{Inst}]; \mathsf{IB} \subseteq \mathsf{OB}$.

Proof. We show that if ib is not already total we can extend it (and maybe ob) into a strictly bigger relation satisfying the constraints of the theorem. Let us assume that there is $(a, b) \in \mathsf{Event}^2$ such that $(a, b) \notin \mathsf{ib}$ and $(b, a) \notin \mathsf{ib}$. We arbitrarily decide to include (a, b) in our relation and we define $\mathsf{ib}' = (\mathsf{ib} \cup \{(a, b)\})^+$ and $\mathsf{ob}' = (\mathsf{ob} \cup [\mathtt{Inst}]; \mathsf{ib}')^+$.

Clearly ib' and ob' are transitive, ib' is irreflexive, and $[\mathtt{Inst}]; \mathsf{ib}' \subseteq \mathsf{ob}'$. We need to prove the following two facts: ob' is still irreflexive; and $\mathsf{ob}'; [\mathtt{Inst}] \subseteq \mathsf{ib}'$.

First, let us check that $(\mathsf{ob} \cup [\mathtt{Inst}]; \mathsf{ib}')^+$ is irreflexive. Since ob and $([\mathtt{Inst}]; \mathsf{ib}')$ are both transitive and irreflexive, a cycle would only be possible by alternating between the two components, so it is enough to show that $(\mathsf{ob}; ([\mathtt{Inst}]; \mathsf{ib}'))^+$ is irreflexive. But $(\mathsf{ob}; ([\mathtt{Inst}]; \mathsf{ib}'))^+ = ((\mathsf{ob}; [\mathtt{Inst}]); \mathsf{ib}')^+ \subseteq (\mathsf{ib}; \mathsf{ib}')^+ \subseteq \mathsf{ib}'$ is irreflexive. Thus ob' is irreflexive.

Then, we need to check that $\mathsf{ob}'; [\mathtt{Inst}] \subseteq \mathsf{ib}'$. Using some rewriting, $\mathsf{ob}' = (\mathsf{ob} \cup [\mathtt{Inst}]; \mathsf{ib}')^+ = \mathsf{ob} \cup (\mathsf{ob}^*; ([\mathtt{Inst}]; \mathsf{ib}'))^+; \mathsf{ob}^*$. We know $\mathsf{ob}; [\mathtt{Inst}] \subseteq \mathsf{ib}'$, which also implies $\mathsf{ob}^*; [\mathtt{Inst}] \subseteq \mathsf{ib}'^*$. So $\mathsf{ob}'; [\mathtt{Inst}] = \mathsf{ob}; [\mathtt{Inst}] \cup ((\mathsf{ob}^*; [\mathtt{Inst}]); \mathsf{ib}')^+; (\mathsf{ob}^*; [\mathtt{Inst}]) \subseteq \mathsf{ib}' \cup (\mathsf{ib}'^*; \mathsf{ib}')^+; \mathsf{ib}'^* \subseteq \mathsf{ib}'$.

Once ib is a total relation on Event, we can similarly freely extend ob into a total relation. □

We use Theorem D.6 above to extend ib and ob into total relations IB and OB.

We use Event to generate new events and some annotated labels. The extended set of events is noted E^{ext}.

- For every $r \in \mathtt{lR}$, from well-formedness conditions, there is w such that $(w, r) \in \mathsf{rf}$ and $\mathsf{eq}_{\mathtt{loc\&v}}(r, w)$. We create an annotated label $\mathtt{lR}\langle r, w\rangle$.
- For every $u \in \mathtt{CAS}$, from well-formedness conditions, there is w such that $(w, u) \in \mathsf{rf}$ and $\mathsf{eq}_{\mathtt{loc\&v}}(u, w)$. We create an annotated label $\mathtt{CAS}\langle u, w\rangle$.
- For every $w \in \mathtt{lW}$ (that is not an initialisation event), we create annotated labels $\mathtt{lW}\langle w\rangle$ and $\mathtt{B}\langle w\rangle$.
- For every $f \in \mathtt{F}$, we create annotated labels $\mathtt{F}\langle f\rangle$.
- For every pair $r \in \mathtt{nlR}$ and $w \in \mathtt{nrW}$ such that $(r, w) \in \mathsf{po}|_{\text{imm}}$, we create two events $a \in \mathsf{Put}$ and $e \in \mathtt{nrEX}$, and the annotated labels $\mathtt{Push}\langle a\rangle$, $\mathtt{NIC}\langle a\rangle$, $\mathtt{nlR}\langle r, w', a, w\rangle$ (where $(w', r) \in \mathsf{rf}$), $\mathtt{nrW}\langle w, e\rangle$, $\mathtt{B}\langle w\rangle$, and $\mathtt{CN}\langle e\rangle$. If there is p such that $(w, p) \in \mathsf{pf}$, we also create an annotated label $\mathtt{P}\langle p, e\rangle$. To simplify later definition, we also extend po such that the event a is placed just before r, and e just after w. I.e., let $\mathsf{po}' = \mathsf{po} \cup \{(e', a) \mid (e', r) \in \mathsf{po}\} \cup \{(a, e') \mid (r, e') \in \mathsf{po}^*\}$ and redefine $\mathsf{po} = \mathsf{po}' \cup \{(e', e) \mid (e, w) \in \mathsf{po}'^*\} \cup \{(e, e') \mid (w, e) \in \mathsf{po}'\}$.
  Note: from well-formedness conditions, every $\mathtt{nlR}$ and every $\mathtt{nrW}$ are part of such a pair.
- Similarly for $\mathtt{nrR}/\mathtt{nlW}$, we create $a \in \mathtt{Get}$, $e \in \mathtt{nlEX}$, $\mathtt{Push}\langle a\rangle$, $\mathtt{NIC}\langle a\rangle$, $\mathtt{nrR}\langle \ldots\rangle$, $\mathtt{nlW}\langle \ldots\rangle$, $\mathtt{B}\langle \ldots\rangle$, and potentially $\mathtt{P}\langle \ldots\rangle$.
- For every $f \in \mathtt{nF}$, we create the annotated labels $\mathtt{Push}\langle f\rangle$, $\mathtt{NIC}\langle f\rangle$, and $\mathtt{nF}\langle f\rangle$.

Then, we use IB and OB to reconstruct a partial path from these annotated labels. We define a path π_0 such that:

- $\pi_0 \in (\mathsf{ALabel} \setminus (\mathsf{Push} \cup \mathtt{NIC} \cup \mathsf{CN}))^*$
- $\mathsf{getI}\lambda(e_1, \pi_0) \prec_{\pi_0} \mathsf{getI}\lambda(e_2, \pi_0) \iff (e_1, e_2) \in \mathsf{IB}$
- $\mathsf{getO}\lambda(e_1, \pi_0) \prec_{\pi_0} \mathsf{getO}\lambda(e_2, \pi_0) \iff (e_1, e_2) \in \mathsf{OB}$

- $\forall w, w' \in \{\mathtt{lW}, \mathtt{nlW}, \mathtt{nrW}\}$, $\mathsf{getI}\lambda(w, \pi_0) \prec_{\pi_0} \mathsf{getO}\lambda(w', \pi_0)$

This is possible from the properties of IB and OB. For pairs of annotated labels not ordered by IB or OB, we decide to order $\mathtt{lW}\langle w\rangle/\mathtt{nlW}\langle w, \_\rangle/\mathtt{nrW}\langle w, \_\rangle$ first and $\mathsf{B}\langle w\rangle$ last. Note that the annotated labels $\mathsf{Push}\langle\ldots\rangle$, $\mathtt{NIC}\langle\ldots\rangle$, and $\mathsf{CN}\langle\ldots\rangle$ not covered by IB/OB are not yet integrated in π_0.

Then we extend π_0 to add annotated labels not considered by the declarative semantics. We use the following extension function that introduces a new annotated label as early as possible after a set of dependencies.

$$\mathsf{extend}(\pi, \lambda, S) \triangleq \begin{cases} \pi_2 \cdot \lambda \cdot \lambda' \cdot \pi_1 & \text{if } \pi = \pi_2 \cdot \lambda' \cdot \pi_1 \wedge \lambda' \in S \wedge \pi_2 \cap S = \emptyset \\ \pi \cdot \lambda & \text{if } \pi \cap S = \emptyset \end{cases}$$

We define a new function to recover the first annotated label corresponding to an event:

$$E^{\mathrm{ext}} \triangleq \mathsf{Event} \cup (\mathtt{Get} \cup \mathtt{Put} \cup \mathtt{nlEX} \cup \mathtt{nrEX})$$

$$\mathsf{getCPU} : E^{\mathrm{ext}} \rightharpoonup \mathsf{ALabel}$$

$$\mathsf{getCPU}(e) \triangleq \begin{cases} \mathsf{getI}\lambda(e, \pi_0) & \text{if } e \in E^{\mathrm{cpu}} = \{\mathtt{lR}, \mathtt{lW}, \mathtt{CAS}, \mathtt{F}, \mathtt{P}\} \\ \mathsf{Push}\langle e\rangle & \text{if } e \in \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\} \\ \text{undefined} & \text{otherwise} \end{cases}$$

And a similar function for events emptying a CPU buffer:

$$\mathsf{getTSO} : E^{\mathrm{ext}} \rightharpoonup \mathsf{ALabel}$$

$$\mathsf{getTSO}(e) \triangleq \begin{cases} \mathsf{B}\langle e\rangle & \text{if } e \in \mathtt{lW} \\ \mathtt{NIC}\langle e\rangle & \text{if } e \in \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\} \\ \text{undefined} & \text{otherwise} \end{cases}$$

Let us consider $(a_1, \ldots, a_n) = \mathsf{Event} \cap \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\}$ in po order, i.e., if $i < j$ then $(a_j, a_i) \notin \mathsf{po}$. We extend π_0 successively until we get π_n:

- We introduce Push as early as possible:
  Let $\pi' = \mathsf{extend}(\pi_{i-1}, \mathsf{Push}\langle a_i\rangle, \{\mathsf{getCPU}(e) \mid (e, a_i) \in \mathsf{po}\})$
- We introduce NIC as early as possible:
  Let $\pi'' = \mathsf{extend}(\pi', \mathtt{NIC}\langle a_i\rangle, \{\mathsf{Push}\langle a_i\rangle\} \cup \{\mathsf{getTSO}(e) \mid (e, a_i) \in \mathsf{po}\})$
- If $a_i \in \mathtt{Put}$, there is $e_i \in \mathtt{nrEX}$ such that $\mathtt{nlR}\langle \_, \_, a_i, w\rangle \prec_{\pi_0} \mathtt{nrW}\langle w, e_i\rangle$. We also introduce CN: Let $S = \{\mathtt{nrW}\langle w, e_i\rangle\} \cup \{\mathtt{nlW}\langle \_, e\rangle \mid (e, e_i) \in \mathsf{po} \cap \mathsf{sqp}\} \cup \{\mathsf{CN}\langle e\rangle \mid (e, e_i) \in \mathsf{po} \cap \mathsf{sqp}\}$, we pose $\pi_i = \mathsf{extend}(\pi'', \mathsf{CN}\langle e_i\rangle, S)$.
  Otherwise, i.e. $a_i \notin \mathtt{Put}$, we simply have $\pi_i = \pi''$

Finally, $\pi = \pi_n$ is our path for an annotated semantics reduction. We clearly have $\mathsf{complete}(\pi)$ by definition. Our goal is then to prove that $\mathsf{wfp}(\pi)$ holds. It is composed of six properties. Note that we already have the existence of the relevant annotated labels, and we need to show that the ordering constraints are respected.

nodup

$\mathsf{nodup}(\pi)$ directly comes from the definition of annotated labels. There is no conflict in event usage.

backComp

Here are a couple lemmas showing that the new annotated labels are not placed too late and do not disturb the expected ordering.

Lemma D.7. *For all $a \in \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\}$ and $b \in \mathsf{Event}$, if $(a, b) \in \mathsf{po}^*$, then $\mathsf{Push}\langle a\rangle \prec_\pi \mathsf{getI}\lambda(b, \pi_0)$.*

Proof. We take an arbitrary b, and proceed for a in po order, i.e., we can assume it holds for $e \in \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\}$ such that $(e, a) \in \mathsf{po}$. By definition, $\mathsf{Push}\langle a\rangle$ comes from an extension $\pi'' = \mathrm{extend}(\pi', \mathsf{Push}\langle a\rangle, \{\mathrm{getCPU}(e) \mid (e, a) \in \mathsf{po}\})$ and has been placed either first—and the result is trivial—or just after some $\mathrm{getCPU}(e)$ with $(e, a) \in \mathsf{po}$. If $e \in \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\}$, we have $\mathsf{Push}\langle e\rangle \prec_{\pi''} \mathsf{Push}\langle a\rangle \prec_{\pi''} \mathrm{getI}\lambda(b, \pi_0)$ by induction hypothesis. If $e \in E^{\mathrm{cpu}} = \{\mathtt{lR}, \mathtt{lW}, \mathtt{CAS}, \mathtt{F}, \mathtt{P}\}$, we have $\mathrm{getI}\lambda(e, \pi_0) \prec_{\pi''} \mathsf{Push}\langle a\rangle \prec_{\pi''} \mathrm{getI}\lambda(b, \pi_0)$ since $(e, b) \in \mathsf{ippo} \subseteq \mathsf{IB}$. □

Lemma D.8. *$\forall a \in \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\}$, $\forall b \in \{\mathtt{nF}, \mathtt{nrR}, \mathtt{nlR}, \mathtt{lW}\}$, if $(a, b) \in \mathsf{po}^*$, then $\mathsf{NIC}\langle a\rangle \prec_\pi \mathrm{getO}\lambda(b, \pi_0)$.*

Proof. We take an arbitrary $b \in \{\mathtt{nF}, \mathtt{nrR}, \mathtt{nlR}\}$, and proceed for a in po order, i.e., we can assume it holds for $e \in \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\}$ such that $(e, a) \in \mathsf{po}$. By definition, $\mathsf{NIC}\langle a\rangle$ comes from an extension $\pi'' = \mathrm{extend}(\pi', \mathsf{NIC}\langle a\rangle, S)$, with $S = \{\mathsf{Push}\langle a\rangle\} \cup \{\mathrm{getTSO}(e) \mid (e, a) \in \mathsf{po}\}$, and has been placed just after some $\lambda \in S$.

- If $\lambda = \mathsf{Push}\langle a\rangle$, then we have $\lambda \prec_{\pi''} \mathsf{NIC}\langle a\rangle \prec_{\pi''} \mathrm{getO}\lambda(b, \pi_0)$ using Lemma D.7 above, since $\mathrm{getI}\lambda(b, \pi_0) = \mathrm{getO}\lambda(b, \pi_0)$.
- If $\lambda = \mathrm{getTSO}\langle e\rangle = \mathsf{NIC}\langle e\rangle$ for some $e \in \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\}$, then we have $\lambda \prec_{\pi''} \mathsf{NIC}\langle a\rangle \prec_{\pi''} \mathrm{getO}\lambda(b, \pi_0)$ by induction hypothesis.
- If $\lambda = \mathrm{getTSO}\langle e\rangle = \mathsf{B}\langle e\rangle$ for some $e \in \mathtt{lW}$, then we have $\mathsf{B}\langle e\rangle \prec_{\pi''} \mathsf{NIC}\langle a\rangle \prec_{\pi''} \mathrm{getO}\lambda(b, \pi_0)$ since $(e, b) \in \mathsf{oppo} \subseteq \mathsf{OB}$.

□

Lemma D.9. *Forall w,e,p, if $\mathtt{nrW}\langle w, e\rangle \in \pi$ and $\mathsf{P}\langle p, e\rangle \in \pi$, then $\mathsf{CN}\langle e\rangle \prec_\pi \mathsf{P}\langle p, e\rangle$.*

Proof. Once again, we proceed for e in po order, i.e., we can assume the result holds for $e' \in \mathtt{nrEX}$ such that $(e', e) \in \mathsf{po}$. $\mathsf{CN}\langle e\rangle$ is inserted in some operation $\pi'' = \mathrm{extend}(\pi', \mathsf{CN}\langle e\rangle, S)$, with $S = \{\mathtt{nrW}\langle w, e\rangle\} \cup \{\mathtt{nlW}\langle \_, e'\rangle \mid (e', e) \in \mathsf{po} \cap \mathsf{sqp}\} \cup \{\mathsf{CN}\langle e'\rangle \mid (e', e) \in \mathsf{po} \cap \mathsf{sqp}\}$. It is then placed just after some label $\lambda \in S$.

- If $\lambda = \mathtt{nrW}\langle w, e\rangle$, we have $\lambda \prec_{\pi''} \mathsf{CN}\langle e\rangle \prec_{\pi''} \mathsf{P}\langle p, e\rangle$ because $(w, p) \in \mathsf{pf} \subseteq \mathsf{IB}$.
- If $\lambda = \mathsf{CN}\langle e'\rangle$ with $(e', e) \in \mathsf{po} \cap \mathsf{sqp}$, then there is some w' such that $(w', w) \in \mathsf{po} \cap \mathsf{sqp}$ and $\mathtt{nrW}\langle w', e'\rangle \in \pi'$. From well-formedness condition number 5 (see Definition 4.3), there is some p' such that $(w', p') \in \mathsf{pf}$ and $(p', p) \in \mathsf{po}$. By induction hypothesis, we have $\mathsf{CN}\langle e'\rangle \prec_{\pi'} \mathsf{P}\langle p', e'\rangle$, and from $(p', p) \in \mathsf{IB}$ we have $\mathsf{P}\langle p', e'\rangle \prec_{\pi'} \mathsf{P}\langle p, e\rangle$. In the end, we have the result $\mathsf{CN}\langle e'\rangle \prec_{\pi''} \mathsf{CN}\langle e\rangle \prec_{\pi''} \mathsf{P}\langle p, e\rangle$.
- If $\lambda = \mathtt{nlW}\langle w', e'\rangle$ with $(e', e) \in \mathsf{po} \cap \mathsf{sqp}$, then we also have $(w', w) \in \mathsf{po} \cap \mathsf{sqp}$, so from well-formedness condition number 5 (see Definition 4.3), there is some p' such that $(w', p') \in \mathsf{pf}$ and $(p', p) \in \mathsf{po}$. We have $\mathtt{nlW}\langle w', e'\rangle \prec_{\pi''} \mathsf{CN}\langle e\rangle \prec_{\pi''} \mathsf{P}\langle p', e'\rangle \prec_{\pi''} \mathsf{P}\langle p, e\rangle$.

□

We can then check that we have $\mathrm{backComp}(\pi)$:

- $\mathtt{lW}\langle w\rangle \prec_\pi \mathsf{B}\langle w\rangle$ comes from the third property when defining π_0; similarly for $\mathtt{nlW}$ and $\mathtt{nrW}$.
- $\mathsf{Push}\langle a\rangle \prec_\pi \mathsf{NIC}\langle a\rangle$ comes from the extension process.
- $\mathsf{NIC}\langle f\rangle \prec_\pi \mathtt{nF}\langle f\rangle$ comes from Lemma D.8; similarly for $\mathsf{NIC}\langle a\rangle \prec_\pi \mathtt{nlR}/\mathtt{nrR}\langle \ldots\rangle$.
- $\mathtt{nlR}\langle r, w, a, w'\rangle \prec_\pi \mathtt{nrW}\langle w', e\rangle$ comes from $(r, w') \in \mathsf{ippo} \subseteq \mathsf{IB}$; similarly for $\mathtt{nrR}/\mathtt{nlW}$.
- $\mathtt{nrW}\langle w, e\rangle \prec_\pi \mathsf{CN}\langle e\rangle$ comes from the extension process
- $\mathtt{nlW}\langle w, e\rangle \prec_\pi \mathsf{B}\langle w\rangle \prec_\pi \mathsf{P}\langle p, e\rangle$ comes from $(w, p) \in [\mathtt{nlW}]; \mathsf{pf} \subseteq \mathsf{OB}$.
- $\mathsf{CN}\langle e\rangle \prec_\pi \mathsf{P}\langle p, e\rangle$ comes from Lemma D.9.

Thus we have $\mathrm{backComp}(\pi)$.

bufFlushOrd

- $\mathtt{lW}\langle w_1\rangle \prec_\pi \mathtt{lW}\langle w_2\rangle \iff \mathtt{B}\langle w_1\rangle \prec_\pi \mathtt{B}\langle w_2\rangle$ when $t(w_1) = t(w_2)$ comes the fact that $[\mathtt{lW}];\mathsf{po};[\mathtt{lW}] \subseteq (\mathsf{IB} \cup \mathsf{OB})$, so both sides are true if and only if $(w_1, w_2) \in \mathsf{po}$; similarly for nlW and nrW on the same queue pair.
- When $t(a_1) = t(a_2)$, $\mathtt{Push}\langle a_1\rangle \prec_\pi \mathtt{Push}\langle a_2\rangle \iff \mathtt{NIC}\langle a_1\rangle \prec_\pi \mathtt{NIC}\langle a_2\rangle \iff (a_1, a_2) \in \mathsf{po}$ from the definition of the extension process (to define π_n).
- For $a \in \{\mathtt{Put}, \mathtt{Get}, \mathtt{nF}\}$, $w \in \mathtt{lW}$, such that $t(a) = t(w)$:
  - If $(w, a) \in \mathsf{po}$, then $\mathtt{lW}\langle w\rangle \prec_\pi \mathtt{Push}\langle a\rangle$ and $\mathtt{B}\langle w\rangle \prec_\pi \mathtt{NIC}\langle a\rangle$ from the definition of the extension process.
  - If $(a, w) \in \mathsf{po}$, then $\mathtt{Push}\langle a\rangle \prec_\pi \mathtt{lW}\langle w\rangle$ and $\mathtt{NIC}\langle a\rangle \prec_\pi \mathtt{B}\langle w\rangle$ from Lemmas D.7 and D.8.
- When $t(w) = t(f)$, $\mathtt{lW}\langle w\rangle \prec_\pi \mathtt{F}\langle f\rangle$ implies $(w, f) \in \mathsf{po}$ (since $[\mathtt{F}];\mathsf{po};[\mathtt{lW}] \subseteq \mathsf{ippo} \subseteq \mathsf{IB}$), which implies $\mathtt{B}\langle w\rangle \prec_\pi \mathtt{F}\langle f\rangle$ (since $[\mathtt{lW}];\mathsf{po};[\mathtt{F}] \subseteq \mathsf{oppo} \subseteq \mathsf{OB}$); similarly for CAS.
- If $w \in \mathtt{nlW}$, $r \in \mathtt{nlR}$, and $\mathsf{sameqp}(w, r)$, then from well-formedness condition number 8 (see Definition 4.3), either $(w, r) \in \mathsf{nfo}$ or $(r, w) \in \mathsf{nfo}$. If $\mathtt{nlW}\langle w, \_\rangle \prec_\pi \mathtt{nlR}\langle r, \_, \_, \_\rangle$, then $(r, w) \notin \mathsf{nfo}$ (since $\mathsf{nfo} \subseteq \mathsf{IB}$) and $(w, r) \in \mathsf{nfo}$. Thus, $\mathtt{B}\langle w\rangle \prec_\pi \mathtt{nlR}\langle r, \_, \_, \_\rangle$ (since $\mathsf{nfo} \subseteq \mathsf{OB}$); similarly for nrW/nrR.

Thus we have $\mathsf{bufFlushOrd}(\pi)$.

pollOrder

Lemma D.10. *For all* $e_1, e_2 \in \{\mathtt{nlEX}, \mathtt{nrEX}\}$*, such that* $\mathsf{sameqp}(e_1, e_2)$*, let* $\lambda_1 \in \{\mathtt{nlW}\langle \_, e_1\rangle, \mathtt{CN}\langle e_1\rangle\}$*,* $\lambda_2 \in \{\mathtt{nlW}\langle \_, e_2\rangle, \mathtt{CN}\langle e_2\rangle\}$*, then* $(e_1, e_2) \in \mathsf{po} \iff \lambda_1 \prec_\pi \lambda_2$*.*

Proof. By symmetry, we only need to show $(e_1, e_2) \in \mathsf{po} \implies \lambda_1 \prec_\pi \lambda_2$. Once again, we proceed for e_1 in po order, i.e., we can assume the result holds for $e' \in \mathtt{nEX}$ such that $(e', e_1) \in \mathsf{po}$.

- If $\lambda_1 = \mathtt{nlW}\langle w_1, e_1\rangle$ and $\lambda_2 = \mathtt{nlW}\langle w_2, e_2\rangle$, then $(e_1, e_2) \in \mathsf{po}$ implies $(w_1, w_2) \in (\mathsf{po} \cap \mathsf{sqp})$, so $(w_1, w_2) \in \mathsf{ippo} \subseteq \mathsf{IB}$ and $\lambda_1 \prec_\pi \lambda_2$.
- If $\lambda_1 = \mathtt{nlW}\langle w_1, e_1\rangle$ and $\lambda_2 = \mathtt{CN}\langle e_2\rangle$, then by definition of the extension process we have $\lambda_1 \prec_\pi \lambda_2$.
- If $\lambda_1 = \mathtt{CN}\langle e_1\rangle$ and $\lambda_2 = \mathtt{nlW}\langle w_2, e_2\rangle$, then λ_1 is inserted in some operation $\pi'' = \mathsf{extend}(\pi', \mathtt{CN}\langle e_1\rangle, S)$, with $S = \{\mathtt{nrW}\langle \_, e_1\rangle\} \cup \{\mathtt{nlW}\langle \_, e'\rangle \mid (e', e_1) \in \mathsf{po} \cap \mathsf{sqp}\} \cup \{\mathtt{CN}\langle e'\rangle \mid (e', e_1) \in \mathsf{po} \cap \mathsf{sqp}\}$. It is then placed just after some label $\lambda \in S$.
  - If $\lambda = \mathtt{nrW}\langle w_1, e_1\rangle$, we have $\lambda \prec_{\pi''} \lambda_1 \prec_{\pi''} \lambda_2$ because $(w_1, w_2) \in \mathsf{ippo} \subseteq \mathsf{IB}$.
  - If $\lambda = \mathtt{CN}\langle e'\rangle$ or $\lambda = \mathtt{nlW}\langle \_, e'\rangle$, with $(e', e_1) \in \mathsf{po} \cap \mathsf{sqp}$, then by induction hypothesis $\lambda \prec_{\pi''} \lambda_1 \prec_{\pi''} \lambda_2$.
- If $\lambda_1 = \mathtt{CN}\langle e_1\rangle$ and $\lambda_2 = \mathtt{CN}\langle e_2\rangle$, then by definition of the extension process we have $\lambda_1 \prec_\pi \lambda_2$.

□

Let us assume we have $e_1, e_2, p_2, \lambda_1, \lambda_2$ such that $\mathsf{sameqp}(e_1, e_2)$, $\lambda_1 \in \{\mathtt{nlW}\langle \_, e_1\rangle, \mathtt{CN}\langle e_1\rangle\}$, $\lambda_2 \in \{\mathtt{nlW}\langle \_, e_2\rangle, \mathtt{CN}\langle e_2\rangle\}$, $\lambda_1 \prec_\pi \lambda_2$, and $\mathtt{P}\langle p_2, e_2\rangle \in \pi$.

From the creation of the events e_1 and e_2, there is some $w_1, w_2 \in \{\mathtt{nlW}, \mathtt{nrW}\}$ such that $(w_i, e_i) \in \mathsf{po}|_{\mathrm{imm}}$. From Lemma D.10, we have $(e_1, e_2) \in \mathsf{po}$ and thus $(w_1, w_2) \in (\mathsf{po} \cap \mathsf{sqp})$. By definition, we also have $(w_2, p_2) \in \mathsf{pf}$. From well-formedness condition number 5 (see Definition 4.3), there is some p_1 such that $(w_1, p_1) \in \mathsf{pf}$ and $(p_1, p_2) \in \mathsf{po}$. Thus we have $\mathtt{P}\langle p_1, e_1\rangle \prec_\pi \mathtt{P}\langle p_2, e_2\rangle$ as required to prove $\mathsf{pollOrder}(\pi)$.

nicActOrder

Let a_1 and a_2 such that $\mathtt{NIC}\langle a_1\rangle \prec_\pi \mathtt{NIC}\langle a_2\rangle$ and $\mathsf{sameqp}(a_1, a_2)$. From the definition of the extension process, we have $(a_1, a_2) \in \mathsf{po}$.

- If $a_1 \in \mathtt{nF}$ or $a_2 \in \mathtt{nF}$, then most of the required results hold by definition of ippo. The only exception is $\mathsf{CN}\langle e\rangle \prec_\pi \mathtt{nF}\langle a_2\rangle$ which holds (by induction on e in po order) because all the dependencies of $\mathsf{CN}\langle e\rangle$ are before $\mathtt{nF}\langle a_2\rangle$ by ippo.
- If $(a_1 \in \mathtt{Get} \wedge a_2 \in \mathtt{Get})$, the result holds by ippo.
- If $(a_1 \in \mathtt{Get} \wedge a_2 \in \mathtt{Put})$, the result holds by Lemma D.10.
- If $(a_1 \in \mathtt{Put} \wedge a_2 \in \mathtt{Get})$, the first results holds by ippo, the second by Lemma D.10.
- If $(a_1 \in \mathtt{Put} \wedge a_2 \in \mathtt{Put})$, the first two results hold by ippo, the last one by Lemma D.10.

Thus we have $\mathsf{nicActOrder}(\pi)$.

wfrd

Let us assume we have $\pi = \pi_4 \cdot \lambda_r \cdot \pi_3$, with $\lambda_r \in \{\mathtt{lR}\langle r, w\rangle, \mathtt{CAS}\langle r, w\rangle, \mathtt{nlR}\langle r, w, \_, \_\rangle, \mathtt{nrR}\langle r, w, \_, \_\rangle\}$. In all cases we have $(w, r) \in \mathsf{rf}$. Another important fact is that $\forall w', (w, w') \in \mathsf{mo} \implies (r, w') \in \mathsf{rb}$.

- If $\lambda_r = \mathtt{lR}\langle r, w\rangle$, we need to show $\mathsf{wfrdCPU}(r, w, \pi_3)$.
  - If $w = init_{\mathtt{loc}(w)}$, then we need to check that $\{\mathtt{B}\langle w'\rangle, \mathtt{CAS}\langle w', \_\rangle \in \pi_3 \mid \mathtt{loc}(w') = \mathtt{loc}(r)\} = \emptyset$ and $\{\mathtt{lW}\langle w''\rangle \in \pi_3 \mid \mathtt{loc}(w'') = \mathtt{loc}(r) \wedge t(w'') = t(r)\} = \emptyset$. For the first, such a w' would imply $(r, w') \in \mathsf{rb} \subseteq \mathsf{OB}$, which contradicts the ordering with λ_r. For the second, such an w'' would imply $(r, w'') \in \mathsf{rb}_\mathsf{b} \subseteq \mathsf{IB}$, and $\lambda_r \prec_\pi \mathtt{lW}\langle w''\rangle$ which similarly contradicts the ordering with λ_r.
  - If $w \in \mathtt{lW}$, $t(w) = t(r)$, and $\mathtt{B}\langle w\rangle \notin \pi_3$. From $(w, r) \in \mathsf{rf}_\mathsf{b} \subseteq \mathsf{IB}$, we have $\lambda_w = \mathtt{lW}\langle w\rangle \prec_\pi \lambda_r$, i.e., $\pi_3 = \pi_2 \cdot \lambda_w \cdot \pi_1$. We need to show that $\{\mathtt{lW}\langle w'\rangle \in \pi_2 \mid \mathtt{loc}(w') = \mathtt{loc}(r) \wedge t(w') = t(r)\} = \emptyset$. Such a w' would imply $(w, w') \in \mathsf{po}$ (from $[\mathtt{lW}]; \mathsf{po}; [\mathtt{lW}] \subseteq \mathsf{ippo} \subseteq \mathsf{IB}$, and the execution graph forcing either $(w, w') \in \mathsf{po}$ or $(w', w) \in \mathsf{po}$), $(w, w') \in \mathsf{mo}$ (from $[\mathtt{lW}]; \mathsf{po}; [\mathtt{lW}] \subseteq \mathsf{oppo} \subseteq \mathsf{OB}$, and well-formedness conditions forcing either $(w, w') \in \mathsf{mo}$ or $(w', w) \in \mathsf{mo}$), and $(r, w') \in \mathsf{rb}_\mathsf{b} \subseteq \mathsf{IB}$ would contradicts the ordering with λ_r.
  - Else we have $\lambda_w \in \pi_3$, with $\lambda_w \in \{\mathtt{B}\langle w\rangle, \mathtt{CAS}\langle w, \_\rangle\}$. If $w \in \mathtt{lW}$ and $t(w) = t(r)$, this is the remaining, else it comes from $(w, r) \in \mathsf{rf}_{\bar{\mathsf{b}}} \subseteq \mathsf{OB}$. Thus we have $\pi_3 = \pi_2 \cdot \lambda_w \cdot \pi_1$, and we need to check two properties. First, we check that
    $\{\mathtt{B}\langle w'\rangle, \mathtt{CAS}\langle w', \_\rangle \in \pi_2 \mid \mathtt{loc}(w') = \mathtt{loc}(r)\} = \emptyset$. It holds because such a w' would again imply $(r, w') \in \mathsf{rb} \subseteq \mathsf{OB}$ and contradict the ordering with λ_r. Second, we check that $\left\{ w' \;\middle|\; \begin{array}{l} \mathtt{lW}\langle w'\rangle \in \pi_3 \wedge \mathtt{B}\langle w'\rangle \notin \pi_3 \wedge \\ \mathtt{loc}(w') = \mathtt{loc}(r) \wedge t(w') = t(r) \end{array} \right\} = \emptyset$. It holds because such a w' would again imply $(w, w') \in \mathsf{mo}$, $(r, w') \in \mathsf{rb}_\mathsf{b} \subseteq \mathsf{IB}$ and contradict the ordering with λ_r.
- If $\lambda_r = \mathtt{CAS}\langle r, w\rangle$, we similarly check that $\mathsf{wfrdCPU}(r, w, \pi_3)$ holds. The difference is that cases that previously contradicted ($\mathsf{rb}_\mathsf{b} \subseteq \mathsf{IB}$) now contradict $\mathsf{bufFlushOrd}(\pi)$ that forces the buffer of $t(r)$ to be empty when performing λ_r.
- If $\lambda_r = \mathtt{nlR}\langle r, w, \_, \_\rangle$, we need to show $\mathsf{wfrdNIC}(r, w, \pi_3)$.
  - If $w = init_{\mathtt{loc}(w)}$, then we need to check that $\{\mathtt{B}\langle w'\rangle, \mathtt{CAS}\langle w', \_\rangle \in \pi_3 \mid \mathtt{loc}(w') = \mathtt{loc}(r)\} = \emptyset$. Such a w' would imply $(r, w') \in \mathsf{rb} \subseteq \mathsf{OB}$, which contradicts the ordering with λ_r.
  - Else we have $\lambda_w \in \pi_3$, with $\lambda_w \in \{\mathtt{B}\langle w\rangle, \mathtt{CAS}\langle w, \_\rangle\}$. This comes from $(w, r) \in \mathsf{rf}_{\bar{\mathsf{b}}} \subseteq \mathsf{OB}$. Thus we have $\pi_3 = \pi_2 \cdot \lambda_w \cdot \pi_1$, and we need to check that
    $\{\mathtt{B}\langle w'\rangle, \mathtt{CAS}\langle w', \_\rangle \in \pi_2 \mid \mathtt{loc}(w') = \mathtt{loc}(r)\} = \emptyset$. It holds because such a w' would again imply $(r, w') \in \mathsf{rb} \subseteq \mathsf{OB}$ and contradict the ordering with λ_r.
- If $\lambda_r = \mathtt{nrR}\langle r, w, \_, \_\rangle$, we similarly check that $\mathsf{wfrdNIC}(r, w, \pi_3)$ for the same reasons.

Thus we have $\mathsf{wfrd}(\pi)$.

Theorem D.11. *Let G be a well-formed consistent execution graph generated from a program* P. *Let* π *be the path obtained from G by the process defined above. Then there is* M′, QP′ *(such that forall*

$t, \overline{n}$ we have $\mathsf{QP}'(t)(\overline{n}) = \langle \varepsilon, \varepsilon, \mathsf{nEX}^* \rangle$*), and an equivalent path* π' *(producing the same outcome as* π*) such that* $\mathsf{P}, \mathsf{M}_0, \mathsf{B}_0, \mathsf{QP}_0, \varepsilon \Rightarrow^* (\lambda t.\mathtt{skip}), \mathsf{M}', \mathsf{B}_0, \mathsf{QP}', \pi'$.

Proof. From above, we have $\mathsf{wf}(\pi)$. This shows that the program configuration can perform the events described by the annotated labels of π. We need to slightly reorder the path because of failed CAS operations, as the two events (memory fence and local read) might not be immediately after each other in π, and we delay the memory fence operation in π' to match the annotated operational semantics. The remaining part of the proof is simply to check that the command rewritings used when deriving the execution graph from P (see Fig. 23) can be used as $\mathcal{E}$ transitions in the annotated semantics for P, which follows from the definitions. □

D.5 Operational Semantics and Annotated Semantics

We define forgetful functions from annotated configurations to operational configurations. For memories, we replace the write event by the value written. For labels within annotated configurations, we drop some arguments to recover the data structure of the operational semantics.

$$
\begin{aligned}
[\![\cdot]\!]_{\mathrm{M}} &: \mathsf{AMem} \longrightarrow \mathsf{Mem} \\
[\![\mathsf{M}]\!]_{\mathrm{M}} &\triangleq \lambda x. v_{\mathrm{w}}(\mathsf{M}(x))
\end{aligned}
$$

$$
[\![\cdot]\!]_{op} : E^{\text{ext}} \rightharpoonup \left\{ y^{\overline{n}} := x^{n}, y^{\overline{n}} := v, \mathsf{ack}_{\mathsf{p}}, x^{n} := y^{\overline{n}}, x^{n} := v, \mathsf{cn}, \mathtt{rfence}\ \overline{n} \right\}
$$

$$
\begin{array}{rlrl}
[\![\mathtt{lW}(x, v_{\mathrm{w}})]\!]_{op} \triangleq x := v_{\mathrm{w}} & & [\![\mathtt{nrEX}(\overline{n})]\!]_{op} \triangleq \mathsf{ack}_{\mathsf{p}} & \\
[\![\mathtt{nrW}(\overline{y}, v_{\mathrm{r}})]\!]_{op} \triangleq \overline{y} := v_{\mathrm{r}} & & [\![\mathtt{F}]\!]_{op} & \text{is undefined} \\
[\![\mathtt{nlW}(x, v_{\mathrm{w}}, \overline{n})]\!]_{op} \triangleq x := v_{\mathrm{w}} & & [\![\mathtt{P}(\ldots)]\!]_{op} & \text{is undefined} \\
[\![\mathtt{nF}(\overline{n})]\!]_{op} \triangleq \mathtt{rfence}\ \overline{n} & & [\![\mathtt{lR}(\ldots)]\!]_{op} & \text{is undefined} \\
[\![\mathtt{Put}(\overline{y}, x)]\!]_{op} \triangleq \overline{y} := x & & [\![\mathtt{CAS}(\ldots)]\!]_{op} & \text{is undefined} \\
[\![\mathtt{Get}(x, \overline{y})]\!]_{op} \triangleq x := \overline{y} & & [\![\mathtt{nlR}(\ldots)]\!]_{op} & \text{is undefined} \\
[\![\mathtt{nlEX}(\overline{n})]\!]_{op} \triangleq \mathsf{cn} & & [\![\mathtt{nrR}(\ldots)]\!]_{op} & \text{is undefined}
\end{array}
$$

$$
[\![\cdot]\!]_{opl} : E^{\text{ext}} \rightharpoonup \left\{ y^{\overline{n}} := x^{n}, y^{\overline{n}} := v, x^{n} := y^{\overline{n}}, x^{n} := v, \mathsf{cn}, \mathtt{rfence}\ \overline{n} \right\}
$$

$$
[\![l]\!]_{opl} = \begin{cases} \mathsf{cn} & \text{if } l = \mathsf{nrEX}(\overline{n}) \\ [\![l]\!]_{op} & \text{otherwise} \end{cases}
$$

The labels that cannot appear in a well-formed annotated configuration are not mapped. For put operations, the operational semantics uses both $(\mathsf{ack}_{\mathsf{p}})$ and (cn) while the annotated semantics uses the label nrEX, so the mapping is different for labels in $\mathbf{wb}_{\mathrm{L}}$.

$[\![\cdot]\!]_{op}$ and $[\![\cdot]\!]_{opl}$ are extended to lists in an obvious way.

We then extend this to configurations as expected. We overload notations to simplify the formulas.

For $\mathsf{qp} = \langle \mathbf{pipe}, \mathbf{wb}_{\mathrm{R}}, \mathbf{wb}_{\mathrm{L}} \rangle \in \mathsf{AQPair}$, we define $[\![\mathsf{qp}]\!] \triangleq \langle [\![\mathbf{pipe}]\!]_{op}, [\![\mathbf{wb}_{\mathrm{R}}]\!]_{op}, [\![\mathbf{wb}_{\mathrm{L}}]\!]_{opl} \rangle$.

For $\mathsf{QP} \in \mathsf{AQPMap}$, we define $[\![\mathsf{QP}]\!] \triangleq \lambda t. \lambda \overline{n}. [\![\mathsf{QP}(t)(\overline{n})]\!]$.

For $\mathsf{B} \in \mathsf{ASBMap}$, we define $[\![\mathsf{B}]\!] \triangleq \lambda t. [\![\mathsf{B}(t)]\!]_{op}$.

Theorem D.12. *For all* $\mathsf{P}, \mathsf{P}' \in \mathsf{Prog}$, $\mathsf{M}, \mathsf{M}' \in \mathsf{AMem}$, $\mathsf{B}, \mathsf{B}' \in \mathsf{ASBMap}$, $\mathsf{QP}, \mathsf{QP}' \in \mathsf{AQPMap}$, $\pi, \pi' \in \mathsf{Path}$, *if* $\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi \Rightarrow \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}', \pi'$ *and* $\mathrm{wf}(\mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi)$, *then* $\mathsf{P}, [\![\mathsf{M}]\!]_{M}, [\![\mathsf{B}]\!], [\![\mathsf{QP}]\!], \pi \Rightarrow \mathsf{P}', [\![\mathsf{M}']\!]_{M}, [\![\mathsf{B}']\!], [\![\mathsf{QP}']\!], \pi'$.

Proof. By straightforward induction on $\Rightarrow$. □

Theorem D.13. *For all* $\mathsf{M} \in \mathsf{AMem}$, $\mathsf{M}'' \in \mathsf{Mem}$, $\mathsf{B} \in \mathsf{ASBMap}$, $\mathsf{B}'' \in \mathsf{SBMap}$, $\mathsf{QP} \in \mathsf{AQPMap}$, $\mathsf{QP}'' \in \mathsf{QPMap}$, *and* $\pi \in \mathsf{Path}$, *if* $\mathsf{P}, [\![\mathsf{M}]\!]_{M}, [\![\mathsf{B}]\!], [\![\mathsf{QP}]\!] \Rightarrow \mathsf{P}', \mathsf{M}'', \mathsf{B}'', \mathsf{QP}''$ *and* $\mathrm{wf}(\mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi)$, *then there exists* $\mathsf{M}' \in \mathsf{AMem}$, $\mathsf{B}' \in \mathsf{ASBMap}$, $\mathsf{QP}' \in \mathsf{AQPMap}$, *and* $\pi' \in \mathsf{Path}$ *such that* $[\![\mathsf{M}']\!]_{M} = \mathsf{M}''$, $[\![\mathsf{B}']\!] = \mathsf{B}''$, $[\![\mathsf{QP}']\!] = \mathsf{QP}''$, *and* $\mathsf{P}, \mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi \Rightarrow \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}', \pi'$.

Proof. By straightforward induction on $\Rightarrow$. In some cases, the reduction enforces a specific annotated label λ and we have $\pi' = \lambda \cdot \pi$; we then need $\mathrm{wf}(\mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi)$ to check that λ is fresh enough for π. □

Theorem D.14 (Operational and Annotated Semantics Equivalence). *For all program* P.

- $[[\mathsf{M}_0]]_M$, $[[\mathsf{B}_0]]$, *and* $[[\mathsf{QP}_0]]$ *are the initialisation for the operational semantics;*
- *If* $\mathsf{P}, \mathsf{M}_0, \mathsf{B}_0, \mathsf{QP}_0, \varepsilon \Rightarrow^* \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}', \pi'$ *then* $\mathsf{P}, [[\mathsf{M}_0]]_M, [[\mathsf{B}_0]], [[\mathsf{QP}_0]] \Rightarrow^* \mathsf{P}', [[\mathsf{M}']]_M, [[\mathsf{B}']], [[\mathsf{QP}']]$
- *If* $\mathsf{P}, [[\mathsf{M}_0]]_M, [[\mathsf{B}_0]], [[\mathsf{QP}_0]] \Rightarrow^* \mathsf{P}', \mathsf{M}'', \mathsf{B}'', \mathsf{QP}''$ *then there exists* $\mathsf{M}' \in \mathsf{AMem}, \mathsf{B}' \in \mathsf{ASBMap}$, $\mathsf{QP}' \in \mathsf{AQPMap}$, *and* $\pi' \in \mathsf{Path}$ *such that* $[[\mathsf{M}']]_M = \mathsf{M}''$, $[[\mathsf{B}']] = \mathsf{B}''$, $[[\mathsf{QP}']] = \mathsf{QP}''$, *and* $\mathsf{P}, \mathsf{M}_0, \mathsf{B}_0, \mathsf{QP}_0, \varepsilon \Rightarrow^* \mathsf{P}', \mathsf{M}', \mathsf{B}', \mathsf{QP}', \pi'$.

Proof. The first point comes from unfolding the definitions. The other two are proved by straightforward induction on $\Rightarrow^*$ and using Theorems D.12 and D.13. The condition $\mathsf{wf}(\mathsf{M}, \mathsf{B}, \mathsf{QP}, \pi)$ is obtained by applying Theorem D.2 when needed. □