Capturing Register and Control Dependence in Memory Consistency Models with Applications to the Itanium Architecture

Lisa Higham[(1)] LillAnne Jackson[(1)(2)] Jalal Kawash[(1)(3)]
higham@cpsc.ucalgary.ca jackson@cpsc.ucalgary.ca kawash@cpsc.ucalgary.ca

[(1)]Department of Computer Science The University of Calgary, Calgary, Canada
[(2)]Department of Computer Science The University of Victoria, Victoria, Canada
[(3)]Department of Computer Science American University of Sharjah, UAE

Abstract

A *complete* framework for modelling memory consistency that includes register and control dependencies is presented. It allows us to determine whether or not a given computation could have arisen from a given program running on a given multiprocessor architecture. The framework is used to provide an exact description of the computations of (a subset of) the Itanium instruction set on an Itanium multiprocessor architecture. We show that capturing register and control dependencies is crucial: a producer/consumer problem is solvable without using strong synchronization primitives on Itanium multiprocessors, but is impossible without exploiting these dependencies.

Keywords: Multiprocessor Memory consistency, register and control dependency, Itanium, process coordination.

1 Introduction

To overcome inefficiency bottlenecks, modern shared memory multiprocessors have complicated memory organization such as replication in caches and write buffers, multiple buses, and out-of-order memory accesses. This causes processors to have differing views of the shared memory, which satisfy only some weak consistency guarantees. To program such systems, programmers need a precise specification of these guarantees, expressed as constraints on the outcomes of executions of programs. We provide a general and intuitive framework for defining the complete memory consistency model of a multiprocessor architecture including its control and register dependencies. The framework facilitates the specification of the exact set of computations that can arise from a multiprocessor system.

As a running example, we illustrate our framework with the Itanium programming language and architecture. Our techniques are exploited to specify the sets of computations that can arise from a multiprocessor program that uses (a subset of) the Intel Itanium instruction set when it is executed on an Itanium multiprocessor architecture. Since Itanium multiprocessors have a complicated architecture for which a complete correct description of its memory consistency does not exist in the literature, this is a second contribution of this paper. Itanium is also to show that including register and control dependencies in any memory consistency model is crucial. We prove that a certain producer/consumer coordination problem is solvable without using strong synchronization primitives on Itanium multiprocessors, but is impossible without exploiting these dependencies.

Our complete framework is an extension of our previous work [5]. It maintains the ability to describe systems at different levels of abstraction and to prove their equivalence. But our previous framework could only capture constraints that arise from shared memory; it did not capture the constraints that arise from control dependencies in an individual processor's program, nor consistency constraints due to private register

dependencies. While our previous framework could be used to determine if a given computation could arise from a given shared memory architecture, it did not extend to answering if the computation could arise from executing a given program on that architecture. These shortcoming are corrected in this paper. We have also used similar ideas to explore the possibility or impossibility of implementing some objects on various weak memory consistency machines [7, 8, 9, 10]. A common feature of all of these is the notion of abstract specifications and implementations, as a generalization of the methodology used for algorithm design in linearizable systems [4]. Even though the systems we consider are far weaker than linearizable ones, we can compose our implementations to achieve implementations of abstract objects on various multiprocessors with weak memory models (or to prove impossibilities).

Local dependencies are modelled in [1, 2]. Indeed, the framework of [1] contains many of the features of our framework, including the association of programs with instantiations of programs, which is similar to our computations, and the capturing of branching and register dependencies. We are unaware, however, of how to use that framework to describe systems at various levels and to prove equivalence.

There are several models for Itanium computations in the literature. Chatterjee and Gopalakrishnan [3] presented an operational model that they describe as a simple formal model for Itanium. Their model is simplified and does not include data dependencies. Yang, Gopalakrishnan, *et. al.* [17] specify Itanium memory ordering rules in terms that can be applied to verification using Prolog with a finite domain constant solver and a boolean Satisfiability checker. Joshi, Lamport, *et. al.* [15] applied the TLA+ specification language to the Itanium manual specifications [13] and use the TLC model checker on the resulting specifications. It was this work that brought about a very clear description of the Itanium memory order [11]. A formal definition of local dependence order, however, is also external to their framework.

In order to focus on the framework description, we make some simplifying assumptions about the Itanium system. For example, we assume that: there is no overlap of memory accesses; we do not consider the effect of store release, load acquire, semaphore or memory fence instructions; we consider all memory to be in the same coherence domain; we assume that all memory locations are cacheable (called WB in the Itanium manuals); and we do not consider procedure stack adjustments. Other work in progress [14] incorporates these features of the Itanium system into our framework.

2 Modelling Multiprocessor Systems and their Computations

We model a multiprocessor system as a collection of *programs* operating on a collection of *objects* under some constraints called a *memory consistency model*.

Informally, as each processor of such a system executes its program, it issues a sequence of *operation invocations* on the objects of the system, and receives *operation responses* for each invocation. Matching each response with its invocation provides each processor with a sequence of *operations*, which is its "view" of the execution. We think of a *computation* as the collection of these views — one for each processor in the system. Because the response of an invocation is determined by the programs of the other processors, the asynchronous interaction of the processors, and the memory organization of the system, there are typically many computations possible for each system. For any such system we seek a way to determine exactly what computations are possible.

Although our framework is general, and can be used for any multiprocessor system, this paper focuses on the computations that can arise from a given program running on some typical multiprocessor architecture, say $\mathcal{A}$, where the objects are private registers and shared memory locations. We use a running example based on the Intel Itanium multiprocessor architecture to illustrate the definitions and their use.

Instructions and programs: The *programming language* of $\mathcal{A}$ consists of a collection of (machine specific) operations, called *instructions*, that perform a variety of load, store, arithmetic and logical functions. Instructions are partitioned into two classes:

- Branch instructions are all instructions that contain a conditional or unconditional branch to a specified label including procedure calls and procedure returns, and
- operational instructions are all other instructions.

1.	loop:	ld8	r1 = [r3]	;load r1 with data whose address is in r3
2.		add	r1 = #1, r1	;add 1 to r1
3.		cmp.ge	p1, p0=#3,r1	;if 3 >= r1 then
4.	(p1)	br	loop	; go back to loop
5.		st8	[r4], r1	;store value of r1 into memory whose address is in r4

Figure 1: An Itanium individual program

Branch instructions transfer program control to a target instruction specified by a label, by operating on the value in the program counter register. Operational instructions move data from one register/memory location to another and/or perform arithmetic or logical functions on the data in registers/memory locations.[1]

Example - Itanium instruction: The Itanium operational instruction "ld8 $r2$=[$r1$]" takes the value from the memory location whose address is stored in $r1$, and places that value in $r2$. In any execution, this instruction has associated with it some value v in $r1$, some value w in the memory location addressed by v and the value u stored in $r2$.

An *individual program* is a finite sequence of instructions from the programming language of $\mathcal{A}$.

Example - Itanium individual program: Figure 1 is an example of an Itanium individual program.

Consider an n processor machine with architecture $\mathcal{A}$. A *multiprogram* for this machine is a collection of n individual programs— one assigned to each processor.

Computations: An instruction that is augmented with *arbitrary* associated values (from the allowed domain), for the registers and memory locations that it accesses is a *completed instruction*. For a conditional branch instruction, "`br`(*cond*) label", the domain of register *cond* is $\{0, 1\}$, and associating one of these values resolves the conditional branch. So the completed form of this instruction is shortened to exactly one of two read operations that access *cond* and return 0 or 1 (namely, 0=read(*cond*) or 1=read(*cond*)). An *individual computation* of an individual processor is *any* sequence of completed instructions. We do not (yet) care what the instructions are or what values are associated with the instructions. A *(multiprocessor) computation* is a set of individual computations, one for each processor. Notice that a computation is defined non-operationally. It is not a single sequence describing an execution of the whole system, nor is there any relationship required between the register and memory values associated with the different completed instructions. Informally, a multiprocessor computation simply records the sequence of operational instructions performed (and therefore completed) by each processor and the values in condition registers that it read along the way.

Example - 2-processor computations: Three 2-processor computations are given in Figures 2, 3 and 4. (The actual values recored in these computation can be bizarre, but they are still computations.)

p_1	p_2
ld8 $r1$ = [$r3$] (ν_{r3}=1284, $\nu_{m(1284)}$=7, $\nu_{r1}\leftarrow$9)	ld8 $r1$ = [$r3$] (ν_{r3}=1027, $\nu_{m(1027)}$=12, $\nu_{r1}\leftarrow$3)
add $r1$ = #1,$r1$ (ν_{r1}=15 $\nu_{r1}\leftarrow$13)	cmp.ge $p1$,$p0$ = #3,$r1$(ν_{r1}=4, $\nu_{p1}\leftarrow$0)
cmp.ge $p1$,$p0$ = #3,$r1$ (ν_{r1}=7, $\nu_{p1}\leftarrow$1)	add $r1$ = #1,$r1$ (ν_{r1}=6 $\nu_{r1}\leftarrow$19)
1=read($p1$)	0=read($p1$)
add $r1$ = #1,$r1$ (ν_{r1}=7 $\nu_{r1}\leftarrow$0)	st8 [$r4$] = $r1$ (ν_{r4}=1104, ν_{r1}=7, $\nu_{m(1104)}\leftarrow$5)
cmp.ge $p1$,$p0$ = #3, $r1$ (ν_{r1}=21, $\nu_{p1}\leftarrow$1)	
0=read($p1$)	
st8 [$r4$] = $r1$ (ν_{r4}=101, ν_{r1}=17, $\nu_{m(101)}\leftarrow$5)	

Figure 2: Computation 1 (arbitrary values are associated with instructions)

Central goal

We aim to construct a comprehensive framework for defining any multiprocessor architecture $\mathcal{A}$, capable of capturing the control and register dependences of $\mathcal{A}$ and the constraints that arise from the memory

[1] We make some simplifying assumptions, for example, that memory locations do not overlap. Also, the instructions of a program are in a format that has been prepared to be processed by an assembler, i.e., some instructions are preceded by labels and the target of a branch is specified by such a label.

p_1	p_2
ld8 $r1$ = [$r3$] (ν_{r3}=1284, $\nu_{m(1284)}$=7, $\nu_{r1}\leftarrow$9)	ld8 $r1$ = [$r3$] (ν_{r3}=1027, $\nu_{m(1027)}$=12, $\nu_{r1}\leftarrow$3)
add $r1$ = #1,$r1$ (ν_{r1}=15 $\nu_{r1}\leftarrow$13)	add $r1$ = #1,$r1$ (ν_{r1}=4 $\nu_{r1}\leftarrow$9)
cmp.ge $p1$,$p0$ = #3,$r1$ (ν_{r1}=7, $\nu_{p1}\leftarrow$1)	cmp.ge $p1$,$p0$ = #3,$r1$(ν_{r1}=6, $\nu_{p1}\leftarrow$0)
1=read($p1$)	0=read($p1$)
add $r1$ = #1,$r1$ (ν_{r1}=7 $\nu_{r1}\leftarrow$0)	st8 [$r4$] = $r1$ (ν_{r4}=1104, ν_{r1}=7, $\nu_{m(1104)}\leftarrow$5)
cmp.ge $p1$,$p0$ = #3, $r1$ (ν_{r1}=21, $\nu_{p1}\leftarrow$1)	
0=read($p1$)	
st8 [$r4$] = $r1$ (ν_{r4}=101, ν_{r1}=17, $\nu_{m(101)}\leftarrow$5)	

Figure 3: Computation 2 (arbitrary values are associated with instructions

p_1	p_2
ld8 $r1$ = [$r3$] (ν_{r3}=1284, $\nu_{m(1284)}$=1, $\nu_{r1}\leftarrow$1)	ld8 $r1$ = [$r3$] (ν_{r3}=1280, $\nu_{m(1280)}$=5, $\nu_{r1}\leftarrow$5)
add $r1$ = #1,$r1$ (ν_{r1}=1 $\nu_{r1}\leftarrow$2)	add $r1$ = #1,$r1$ (ν_{r1}=5 $\nu_{r1}\leftarrow$6)
cmp.ge $p1$,$p0$ = #3,$r1$ (ν_{r1}=2, $\nu_{p1}\leftarrow$1)	cmp.ge $p1$,$p0$ = #3,$r1$(ν_{r1}=6, $\nu_{p1}\leftarrow$0)
1=read($p1$)	0=read($p1$)
add $r1$ = #1,$r1$ (ν_{r1}=2 $\nu_{r1}\leftarrow$3)	st8 [$r4$] = $r1$ (ν_{r4}=1284, ν_{r1}=0, $\nu_{m(1284)}\leftarrow$0)
cmp.ge $p1$,$p0$ = #3, $r1$ (ν_{r1}=3, $\nu_{p1}\leftarrow$0)	
0=read($p1$)	
st8 [$r4$] = $r1$ (ν_{r4}=1280, ν_{r1}=1, $\nu_{m(1280)}\leftarrow$1)	

Figure 4: Computation 3

organization of $\mathcal{A}$. The definition of any such $\mathcal{A}$ must be precise enough to decide the following *Architecture-Computation* decision problem.

Input: a multiprogram P, containing an individual program for each processor of $\mathcal{A}$, and a multiprocessor computation C.

Question: Could C arise from executing P on $\mathcal{A}$?

Example - An instance of the Architecture-Computation Problem: Let $IP = \{prog_1, prog_2\}$ where $prog_1$ and $prog_2$ are both the program in Figure 1. When applied to our running example, the Architecture-Computation problem becomes: "Which of Computations 1, 2, and 3 (in Figures 2, 3 and 4) could arise from running IP on an Itanium multiprocessor?"

Two major steps are required to answer the Architecture-Computation question. The first is to determine if each processor's individual computation (ignoring the associated values) could have arisen from its program. If the answer is yes, then the second step is to determine whether or not the individual program sequences could have interacted in such a way under architecture $\mathcal{A}$ as to produce the multiprocessor computation — that is, the values associated with each of the instructions.

Program graphs and computational forms: To answer the first question, each processor's program is modelled as a directed graph, where (roughly) vertices represent instructions and edges represent the 'next instruction' relation. We will see that any computation of an individual program must correspond to some path through the program graph.

Specifically, for any individual program *prog* associate a directed node-labeled graph called the *program-graph of* prog, denoted $\mathcal{G}(prog)$, as follows.

Nodes of $\mathcal{G}(prog)$ For every operational instruction `inst`, there is a vertex η(`inst`) with label `inst`. For every conditional branch instruction `br` of the form "br(*cond*) label:", there are two vertices, η_0(`br`) with label 0=read(*cond*) and η_1(`br`) with label 1=read(*cond*). Henceforth, the labels of vertices of $\mathcal{G}(prog)$ are referred to as *node-labels* to distinguish them from the labels of instructions in *prog*.

Directed edges of $\mathcal{G}(prog)$ First, for any instruction `inst` in *prog*, define vertices(`inst`), as follows:

- if `inst` is an operational instruction then vertices(`inst`) = $\{\eta$(`inst`)$\}$,
- if `inst` is a conditional branch instruction then vertices(`inst`) = $\{\eta_0$(`inst`), η_1(`inst`)$\}$,

- if `inst` is an unconditional branch instruction, say "br continue-here", then vertices(`inst`) = vertices($\widehat{\texttt{inst}}$) where $\widehat{\texttt{inst}}$ is the instruction with label "continue-here:".

Denote by `p-succ(inst)` the instruction in *prog* that follows `inst`. For each vertex η in $\mathcal{G}(prog)$, there is a directed edge from η to every vertex in `g-succ(`η`)`, where `g-succ(`η`)` is defined as follows:

- for an operational instruction `inst`, `g-succ(`η`(inst))` = vertices(`p-succ(inst)`) and let $\widehat{\texttt{inst}}$ be the instruction with label "label:", Then `g-succ(`η_0`(br))` = vertices(`p-succ(br)`) and `g-succ(`η_1`(br))` = vertices($\widehat{\texttt{inst}}$).

Example - Program graph for an Itanium program: Here is the program graph of the Itanium individual program of Figure 1.

Let `start` be the first instruction in an individual program *prog*. An *individual program sequence for* prog is the sequence of the node-labels on any (possibly non-simple) directed path in $\mathcal{G}(prog)$ that begins with any vertex in vertices(`start`). Let $P = \{prog_1, prog_2, ..., prog_n\}$ be a multiprogram. Then $CF = \{$*prog-seq*$_1$, *prog-seq*$_2$, ..., *prog-seq*$_n\}$ is a *computational form of* P if and only if for $1 \leq i \leq n$, *prog-seq*$_i$ is an individual program sequence for $prog_i$. A computation C *agrees with* CF if for each i the ith sequence of instructions in C, ignoring values of registers and memory locations, is the same as *prog-seq*$_i$. Claim 2.1 follows immediately from these definitions.

Claim 2.1 *A multiprocessor computation can be a computation of a multiprogram P only if it agrees with a computational form of P.*

Example - Computational form for an Itanium 2-processor multiprogram: $CF = \{$*prog-seq*$_1$, *prog-seq*$_2\}$, where *prog-seq*$_1$ and *prog-seq*$_2$ are given in Figure 5, is a computational form of the Itanium multiprogram IP in the running example.

$prog\text{-}seq_1$		$prog\text{-}seq_1$	
ld8	$r1 = [r3]$	ld8	$r1 = [r3]$
add	$r1 = \#1,r1$	add	$r1 = \#1,r1$
cmp.ge	$p1,p0 = \#3,r1$	cmp.ge	$p1,p0 = \#3,r1$
1=read($p1$)		0=read($p1$)	
add	$r1 = \#1,r1$	st8	$[r4] = r1$
cmp.ge	$p1,p0 = \#3,r1$		
0=read($p1$)			
st8	$[r4] = r1$		

Figure 5: A computational form

Notice that a computational form contains only operational instructions; there are no branch instructions. In the program graph each conditional branch instruction is replaced by one node for each of the two possible outcomes. The outgoing edges of these nodes lead to the next instruction that is correct given the value of the condition. Each unconditional branch is replaced only by edges.

Example - Agreement of computations with an Itanium computational form: Both computations in Figures 3 and 4 agree with the computational form for IP in Figure 5 but the computation in Figure 2 does not because the cmp.ge and add instructions of p_2 are in the opposite order from the order of the add and cmp.ge node-labels in the program graph.

From computational forms to computations: Once it is determined that a computation C agrees with a computational form of the given multiprogram, the second question is whether the collection of individual program sequences in the computation form could have interacted in such a way as to produce the values associated with each of C's instructions. To answer this question, the rules of interaction for the machine with architecture $\mathcal{A}$ are modelled as a collection of constraints on various subsets of the instructions of a computation. For each subset there is a specified partial order, for which there must be a total order extension that is *valid*. Furthermore there may be agreement properties required between some of the total order extensions. We show next how to check the validity of sequences of completed instructions.

Validity

Meaning of instructions: To define validity, each entry in a computational form (i.e. each node-label, ηl) is assigned a *meaning*, denoted $\mathcal{M}(\eta l)$, by mapping it to a short program that uses only read, write, arithmetic and logic operations on variables. Specifically, the programming language `Trivial`:

- has two kinds of objects: single-reader/single-writer atomic variables and multi-reader/multi-writer atomic variables, and
- supports five types of operations: read and write operations on atomic variables, arithmetic and logic operations on single-reader/single-writer atomic variables, and "if-conditional-then-operations-else-operations"

Recall that each node-label, ηl, is either an operational instruction or 0=read($cond$), or 1=read($cond$). For $\eta l \in \{0 = \text{read}(cond), 1 = \text{read}(cond)\}$ define $\mathcal{M}(\eta l)$ to be the identity function. If ηl is an operational instruction then $\mathcal{M}(\eta l)$ is defined to be the short `Trivial` program that re-expresses the semantics of the instruction as specified by the instruction manual of $\mathcal{A}$.

$\mathcal{M}$(*ld8 r1=[r3]*)	$\mathcal{M}$(*add r1=#1,[r3]*)	$\mathcal{M}$(*cmp.eq p1,p2=r1,r3*)	$\mathcal{M}$(*st8 [r4]=r5*)

$\nu_1 = \text{read}(r3)$
$\nu_2 = \text{read}(\nu_1)$
$\text{write}(r1, \nu_2)$
	$\nu_1 = \text{read}(r3)$
	$\nu_2 = \text{read}(\nu_1)$
	$\nu_3 = 1 + \nu_2$
	$\text{write}(r1, \nu_3)$
		$\nu_1 = \text{read}(r2)$
		$\nu_2 = \text{read}(r3)$
		if (ν_1==ν_2) then {
		  $\text{write}(p2, 1)$
		  $\text{write}(p1, 0)$ }
		else { $\text{write}(p2, 0)$
		  $\text{write}(p1, 1)$ }
			$\nu_1 = \text{read}(r5)$
			$\nu_2 = \text{read}(r4)$
			$\text{write}(mem(\nu_2), \nu_1)$

(a)	(b)	(c)	(d)

Figure 6: Meaning of four Itanium instructions

Example - Meaning of some Itanium operational instructions: Figure 6 gives four examples. These can be extracted from the Itanium Instruction Set Reference manual [12].

An operation in `Trivial` is *completed* by assigning a value to each variable in the operation. Recall that a completed instruction, `inst`, has values associated with each register or shared memory location that it accesses, and the variables in $\mathcal{M}$(`inst`) are assigned by reading, writing, or performing arithmetic/logic operations on the values in these registers. If the register values in the completed instruction `inst` are used to compute the values of the corresponding variables in $\mathcal{M}$(`inst`), the resulting sequence of `Trivial` operations is called the *derived completed meaning of* `inst`, denoted $\mathcal{DM}$(`inst`). Because the operations are completed, the conditional operations in this sequence can be resolved, so the derived completed meaning of `inst` is reduced to a sequence I of completed read, write and arithmetic/logic operations. Define the *remote derived completed meaning of* `inst`, denoted $\mathcal{RDM}$(`inst`). to be the subsequence of I consisting of only the completed write operations to shared memory locations.

Example - Derived completed meaning: In the instruction ld8 $r1$=[$r3$] in Figure 6(a) values are assigned to ν_1 and ν_2, for example, $\nu_1 = 17$, $\nu_2 = 21$.

Example - Remote derived completed meaning: For the instruction st8 [$r4$]=$r5$ in Figure 6(d), only the write($mem(\nu_2, \nu_1)$ operation together with the value assigned to ν_1 is included in the remote derived completed meaning.

Definition of validity of a sequence of completed instructions: Let $S = s_1, s_2, \ldots, s_k$ be any sequence of completed instructions of a computation. The *computed meaning of* S is the sequence of `Trivial` completed operations formed by concatenating $\mathcal{DM}(s_i)$ for i from 1 to k. For a given individual program p, the *computed meaning of* S *for* p is the sequence of `Trivial` completed operations formed by concatenating as follows. Replace each s_i that is an instruction by p with $\mathcal{DM}(s_i)$; replace each s_i that is an instruction by a processor different from p with $\mathcal{RDM}(s_i)$. Notice that both the computed meaning of S and the computed meaning of S for p are sequences of read, write and arithmetic/logic operations on atomic variables, so it is straightforward to determine if such a sequence is valid. Specifically, it is *valid* if each read returns the value of the last write to the same variable and all arithmetic/logic operations on variables are correct. Finally, a sequence S of completed instructions from a computation is *valid (respectively, valid for p)* if the computed meaning of S (respectively, the computed meaning of S for p) is valid.

Partial orders and memory consistency models

The final task is to capture the rules that govern executions of multiprograms under architecture $\mathcal{A}$ as defining properties of the computations that can be produced. These properties, collectively called a memory consistency model, are different for every multiprocessor machine, but have the same general structure. They are expressed as a collection of partial orders relations on the completed instructions of a computation. Each of these partial orders is required to have a total order extension that

- is valid and
- shares some agreement properties with other partial orders.

Example - The Itanium memory consistency model[2]: Let $I(\mathcal{C})$ be the set of all completed instructions in a computation C. $I(\mathcal{C})|p$ denotes the subset of $I(\mathcal{C})$ in processor p's program sequence; $I(\mathcal{C})|x$ denotes the subset whose meaning contains a read or write operation on (register or shared memory) variable x; and $I(\mathcal{C})|br$ is the subset whose meaning is 0=read(*cond*) or 1=read(*cond*), where cond is a (condition) register. $I(\mathcal{C})|r$ denotes the subset containing only the instruction instances that contain a read operation on a shared memory variable; $I(\mathcal{C})|w$ the subset containing only the instruction instances that contain a write operation on a shared memory variable. The relation $(I(\mathcal{C}), \xrightarrow{prog})$, called *program order*, is the set of all pairs (i, j) of completed instructions that are in the same individual computation of C and such that i precedes j in that sequence. For any partial order relation $(I(\mathcal{C}), \xrightarrow{y})$, the notation $i \xrightarrow{y} j$ is used interchangeably with $(i, j) \in (I(\mathcal{C}), \xrightarrow{y})$.

Define the following partial orders:

- Local dependence order $(I(\mathcal{C}), \xrightarrow{dep_p})$: For $i, j \in I(\mathcal{C})|p$, $i \xrightarrow{dep_p} j$ if $i \xrightarrow{prog} j$ and either

[2]Itanium provides a rich instruction set, which includes semaphore and fence instructions. The definition formulated here ignores acquire, release, and fence instructions. The development and proofs for the general Itanium definition are elsewhere [6].

Register: $i, j \in I(\mathcal{C})|x$, where x is a register, or

Branch: $i \in I(\mathcal{C})|br$.

- Orderable order $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{ord_p})$ for each $p \in P$: $i \xrightarrow{ord_p} j$ if $i, j \in I(\mathcal{C})|p \cup I(\mathcal{C})|w$ and $i \xrightarrow{prog} j$ and $i, j \in I(\mathcal{C})|x$ and $(i \in I(\mathcal{C})|w$ or $j \in I(\mathcal{C})|w)$

Itanium memory consistency definition: A computation C satisfies Itanium consistency if for each $p \in P$, there is a total order $\xrightarrow{S_p}$ of the operations $I(\mathcal{C})|p \cup I(\mathcal{C})|w$ that is valid for p, such that

1. $(I(\mathcal{C})|p, \xrightarrow{dep_p}) \subseteq (I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{S_p})$, (Local requirement) and
2. $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{ord_p}) \subseteq (I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{S_p})$, (Orderable requirement) and
3. If $i_1, i_2 \in I(\mathcal{C})|x|w$ and $i_1 \xrightarrow{S_p} i_2$ then $i_1 \xrightarrow{S_q} i_2$, $\forall q \in P$, (Same Memory agreement) and
4. There does not exist a cycle of $i_1, i_2 \ldots i_k \in I(\mathcal{C})|w$ where $i_j \in I(\mathcal{C})|p_j, \forall j \in \{1, 2, \ldots k\}$ and $k \leq n$ such that: $i_k \xrightarrow{S_1} i_1$, and $i_1 \xrightarrow{S_2} i_2$, and $i_2 \xrightarrow{S_3} i_3 \ldots$ and $i_{k-1} \xrightarrow{S_k} i_k$ (cycle free agreement)

Define Itanium-Dep (Itanium minus dependence) to be identical to the preceding definition, without the Local requirement.

Example - Itanium-Dep and Itanium consistency: Computation 3 of Figure 4 satisfies Itanium-Dep. The following sequences show that Computation 3 satisfy Itanium-Dep requirements and agreement properties:

S_{p_1} : ld8 $r1$ = [$r3$] (ν_{r3}=1284, $\nu_{m(1284)}$=1, ν_{r1}←1) $\xrightarrow{S_{p_1}}$ st8$_{p_2}$ [$r4_{p_2}$] = $r1_{p_2}$ ($\nu_{r4_{p_2}}$=1284, $\nu_{r1_{p_2}}$=0, $\nu_{m(1284)}$←0) $\xrightarrow{S_{p_1}}$ st8 [$r4$] = $r1$ (ν_{r4}=1280, ν_{r1}=1, $\nu_{m(1280)}$←1) $\xrightarrow{S_{p_1}}$ add $r1$ = #1,$r1$ (ν_{r1}=1 ν_{r1}←2) $\xrightarrow{S_{p_1}}$ cmp.ge $p1$,$p0$ = #3,$r1$ (ν_{r1}=2, ν_{p1}←1) $\xrightarrow{S_{p_1}}$ 1=read($p1$) $\xrightarrow{S_{p_1}}$ add $r1$ = #1,$r1$ (ν_{r1}=2 ν_{r1}←3) $\xrightarrow{S_{p_1}}$ cmp.ge $p1$,$p0$ = #3, $r1$ (ν_{r1}=3, ν_{p1}←0) $\xrightarrow{S_{p_1}}$ 0=read($p1$)

S_{p_2} : st8 [$r4$] = $r1$ (ν_{r4}=1284, ν_{r1}=0, $\nu_{m(1284)}$←0) $\xrightarrow{S_{p_2}}$ ld8 $r1$ = [$r3$] (ν_{r3}=1280, $\nu_{m(1280)}$=5, ν_{r1}←5) $\xrightarrow{S_{p_2}}$ add $r1$ = #1,$r1$ (ν_{r1}=5 ν_{r1}←6) $\xrightarrow{S_{p_2}}$ cmp.ge $p1$,$p0$ = #3, $r1$(ν_{r1}=6, ν_{p1}←0) $\xrightarrow{S_{p_2}}$ 0=read($p1$) $\xrightarrow{S_{p_2}}$ st8$_{p_1}$ [$r4_{p_1}$] = $r1_{p_1}$ (ν_{r4}=1280, $\nu_{r1_{p_1}}$= 1, $\nu_{m(1280)}$←1)

To show that these sequences are valid we provide `Trivial` instructions for the derived completed meaning or remote derived completed meaning for each instruction in the sequence above. We use [ ] to delineate the operations for each instruction.

Validity Sequence for S_{p_1}: [$\nu_1(1284)$ = read($r3$), $\nu_2(1)$ = read(ν_1), write($r1, \nu_2(1)$)] [write($\nu_2(1284), \nu_1(0)$)] [$\nu_1(1)$ = read($r1$), $\nu_2(1280)$ = read($r4$), write($\nu_2(1280), \nu_1(1)$)] [$\nu_1(1)$ = read($r1$), $\nu_2(2) = 1 + \nu_1(1)$, write($r1, \nu_2(2)$)] [$\nu_1(2)$ = read($r1$), if ($\nu_1(2)$ >= #3) then write($p1, 1$) else write($p1, 0$)] [1=read($p1$)] [$\nu_1(2)$ = read($r1$), $\nu_2(3) = 1 + \nu_1(2)$, write($r1, \nu_2(3)$)] [$\nu_1(3)$ = read($r1$), if ($\nu_1(1)$ >= #3) then write($p1, 1$) else write($p1, 0$)] [0=read($p1$)]

Validity Sequence for S_{p_2}: [$\nu_1(0)$ = read($r1$), $\nu_2(1284)$ = read($r4$), write($\nu_2(1284), \nu_1(0)$)] [$\nu_1(1280)$ = read($r3$), $\nu_2(5)$ = read(ν_1), write($r1, \nu_2(5)$)] [$\nu_1(5)$ = read($r1$), $\nu_2(6) = 1 + \nu_1(5)$, write($r1, \nu_2(6)$)] [$\nu_1(6)$ = read($r1$), if ($\nu_1(6)$ >= #3) then write($p1, 1$) else write($p1, 0$)] [0=read($p1$)] [write($\nu_2(1280), \nu_1(1)$)]

However, Computation 3 does not satisfy Itanium because it must extend Local dependence order, which requires that the st8 instruction in each sequence follows the add instructions in that sequence, and this ensures the st8 instructions are invalid. A computation that is similar to Computation 3 except that the st8 instruction by p_1 has values (ν_{r4}=1280, ν_{r1}=3, $\nu_{m(1280)}$←3) and the st8 instruction by p_2 has values (ν_{r4}=1280, ν_{r1}=6, $\nu_{m(1280)}$←6) satisfies Itanium.

3 Consequences of Ignoring Dependence Order

Ignoring register and control dependences can have substantial implications on the underlying memory consistency model. This in turn can lead to erroneous conclusions about the capabilities or limitations of the

architecture under consideration. This section illustrates such an error using a simple producer-consumer example and the Itanium architecture.

3.1 A simple producer-consumer multiprocessor system

Informally, the simple producer-consumer (SPC) multiprocessor system defined here has one producer program, producing items, which are consumed by a single consumer program. The producer and consumer take turns in producing and consuming items. This specialized system can be defined using the framework of Section 2 as follows:[3]

Objects: A *SPC object* X supports two operations P-write(X,ι) and ι=C-read(X). The semantics of these operations are exactly the same as the read and write operations on an atomic variable. However, the validity requirement is substantially stronger.

Validity: A sequence of (C-read and P-write) operations on SPC object X is valid if: (1) it starts with a P-write operation, (2) alternates between C-read and P-write operations, (3) each C-read returns the value written by the immediately preceding P-write operation, and (4) it has an equal number of P-write and C-read operations.[4]

Programs: There are two programs: the producer repeatedly 'P-writes' the SPC object, and the consumer repeatedly 'C-reads' it, for a specified (possibly infinite) number of rounds.

SPC System: The SPC multiprocessor system consists of one SPC object X and a multiprogram of one producer and one consumer. The memory consistency model is sequential consistency[16].

3.2 Proving relationships between multiprocessor systems

Given a multiprocessor system with multiprogram P, objects J, and a memory consistency model M, we call (P, J) a *multiprocess* and when J consists entirely of atomic variables and registers, we call (J, M) an *M platform*.

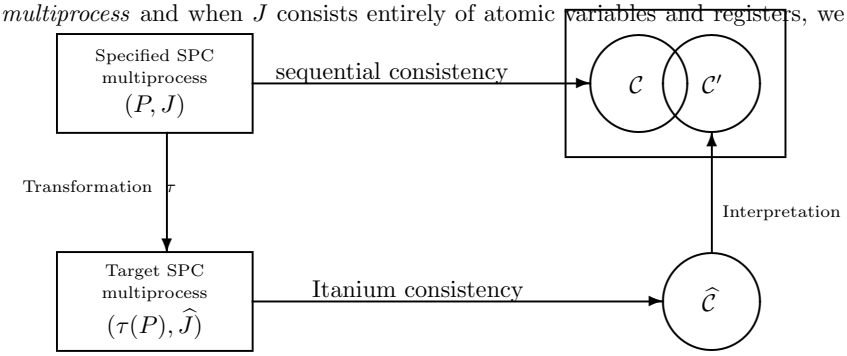

Figure 7: Program transformation and computation interpretation.
If $\mathcal{C}' \subseteq \mathcal{C}$, then the transformation is an implementation.

The *specified* SPC multiprocess (P, J) gives rise to a set of allowable (specification) computations, $\mathcal{C}$ in Figure 7, under sequential consistency. We model the execution of the *specified* multiprocess on the *target* Itanium architecture, as a transformation, τ. This transformation replaces the specified objects, operations, and memory consistency model respectively with target (Itanium) objects, instructions, and memory consistency model. Specifically, an operation on a specified object is *transformed* to the target objects by providing a subroutine for the operation's invocation where this subroutine uses only instruction

[3] A general producer-consumer system definition may allow several producers and consumers and more complex shared objects such as queues, where producers enqueue items and consumers dequeue them. This very restricted definition (two programs and a queue of size one) suffices for this section. Furthermore, without exploiting additional Itanium synchronization mechanisms, such as fences, acquires, or releases, a solution for the general case is not possible.

[4] Note that strict alternation between P-write and C-read operations is not necessary at lower levels. This specification-level requirement can be satisfied at an implementation-level as long as operations appear as if they do alternate.

invocations on the target objects, $\widehat{J}$. If the specified operation returns an output, then the subroutine must return a value of the same type as this output. An object in J is *transformed* to the target object(s) in $\widehat{J}$ by transforming each of its operations. A transformation of each of the objects of a specified multiprocess to objects of the target multiprocess can be naturally extended to a *program transformation* by replacing each operation invocation in the specified multiprogram with the subroutine for that operation invocation. The transformed multiprogram together with the target objects $(\tau(P), \widehat{J})$ comprise the *transformed multiprocess.*

Hence, each P-write and C-read operations will be transformed to subroutines that only make use of Itanium atomic read and write variables and registers. The transformed multiprocess gives rise under Itanium consistency to an allowable set of computations, $\widehat{\mathcal{C}}$. Any computation in $\widehat{\mathcal{C}}$ can be *interpreted* as a computation of the specified multiprogram by attaching to each operation invocation of the specified multiprogram the value returned by the corresponding subroutine. Thus, the set of computations of the target system provides a set of interpreted computations $\mathcal{C}'$ of the specified multiprocess.

A transformation of a specified multiprocess is called an *implementation* of the specified system, if, informally, (in Figure 7) the set $\mathcal{C}'$ of computations produced by traveling the long way around is a non-empty subset of the computations $\mathcal{C}$ allowed by the specified multiprocessor system. The formal details, including stronger (than 'implementations') relations between systems, are elsewhere [5].

3.3 In the presence of dependence

Use p and c to refer to the producer and consumer programs, respectively. For the transformation of the SPC multiprocess given in Figure 8, consider the last load in two consecutive iterations of p. We use dependence order and validity to prove that at least one copy of c's store at line 10 occurs in p's sequence $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{S_p})$ between these two loads. Also consider the last load in two consecutive iterations of c. We use dependence order and validity to prove that at least one copy of p's store at line 4 occurs in c's sequence $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{S_c})$ between these loads. Orderable order and the same memory agreement property force the stores to strictly alternate in each sequence, ensuring that the c loads every produced value.

Produce

Pseudo Code	$\tau((P{-}write(\mathrm{X}, item))$:		
	0	mov r2=Q	; r2 contains the address of Q
while (Q $\neq$ 0) wait;	1 L:	ld8 r1=[r2]	; load r1 with Q
	2	cmp.neq p1,p0 = r1,r0	; is r1 == 0?
	3	(p1) br L:	; If not, go back to L
Q $\longleftarrow$ item	4	st8 [r2]= item	; store item into Q

Consume

Pseudo Code	$\tau((C - read(\mathrm{X}))$ returns item in $r3$:		
	5	mov r2=Q	; r2 contains the address of Q
while (Q = 0) wait;	6 M:	ld8 r1=[r2]	; load r1 with Q
	7	cmp.eq p1,p0 = r1, r0	; is r1 == 0?
	8	(p1) br M:	; If so, go back to M
item $\longleftarrow$ Q	9	ld8 r1=[r2]	; load r1 with Q
Q $\longleftarrow$ 0	10	st8 [r2]=r0	; store 0 into Q
	11	mov r3, r1	; put consumed item into r3

Figure 8: A transformation of the specified SPC object, X, on Itanium, using the target objects $\{Q, r1_p, r2_p, p1_p, r1_c, r2_c, r3_c, p1_c\}$

Theorem 3.1 *Transformation τ in Figure 8 is an implementation of the SPC system on an Itanium platform.*

Proof: The pseudo code in Figure 8 is provided to assist the reader but it is not sufficient to prove correctness of the solution. The program for the transformed producer, p, is the sequence $\tau((P - write(X, item))_1$ $\tau((P - write(X, item))_2$ $\tau((P - write(\mathrm{X}, item))_3$ . . . Similarly the program for the transformed consumer, c, is the sequence $\tau((C - read(X)))_1$ $\tau((C - read(X)))_2$ $\tau((C - read(X)))_3$ . . .

The $f$orm of the program sequences for p and c are given in Figure 9 where base numbers indicated the instructions as labeled in the transformations. Subscripts indicate from which transformation

of $\tau((P - write(X, item))_i$ or $\tau((C - read(X)))_i$ the instruction originates and superscripts distinguish between multiple copies of instructions in the spin loops (instructions 1 through 3 and 6 through 8). This is

Figure 9: Form of Program Sequences for programs p and c

only the form of the program sequences because we have not specified the number of times that spin loops are instantiated in the program sequence. Dotted arrows in Figure 9 show the Local dependence order between instructions in each program sequence. To simply the drawing, transitivity arrows have not been included. In Figure 10 the Orderable order order between the store instructions of each sequence is shown.

Figure 10: Orderable instructions from programs p, c

Recall that each instruction has a meaning that is a program of operations in the simple programming language. In particular, each instruction in the two transformations are mapped to the operations specified in the following tables:

Instruction, in		$\mathcal{M}$(in):
0	mov r2=Q	write$(r1, Q)$
1	ld8 r1=[r2]	ν_{p1}=read$(r2)$, ν_{p2}=read(ν_{p1}), write$(r1, \nu_{p2})$
2	cmp.neq p1,p0 = r1,r0	ν_{p3}=read$(r1)$, if $(\nu_{p3} \neq 0)$ write$(p1, 1)$ else write$(p1, 0)$
3[1]	1 = read$(p1)$	1 = read$(p1)$
3[0]	0 = read$(p1)$	0 = read$(p1)$
4	st8 [r2]= item	ν_{p4}=read$(r2)$, write$(\nu_{p4}, item)$
Instruction, in		$\mathcal{M}$(in):
5	mov r2=Q	write$(r1, Q)$
6	ld8 r1=[r2]	ν_{c1}=read$(r2)$, ν_{c2}=read(ν_{c1}), write$(r1, \nu_{c2})$
7	cmp.eq p1,p0 = r1, r0	ν_{c3}=read$(r1)$, if $(\nu_{c3} = 0)$ write$(p1, 1)$ else write$(p1, 0)$
8[1]	1 = read$(p1)$	1 = read$(p1)$
8[0]	0 = read$(p1)$	0 = read$(p1)$
9	ld8 r1=[r2]	ν_{c4}=read$(r2)$, ν_{c5}=read(ν_{c4}), write$(r1, \nu_{c4})$
10	st8 [r2]=r0	ν_{c5}=read$(r2)$, write$(\nu_{c5}, 0)$
11	mov r3, r1	ν_{c6}=read$(r1)$, write$(r3, \nu_{c6})$

The computation form for these program sequences will include all the instructions in each program sequence, including a decision on the number of instantiations of the spin loops. An Itanium computation, C, supplies values for ν_{pi} and ν_{cj} in a manner that ensures validity of the sequences $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{p})$ and $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ and maintains Local dependence order $(I(\mathcal{C})|p, \xrightarrow{dep_p})$ order and $(I(\mathcal{C})|c, \xrightarrow{dep_p})$ order and weak Orderable order $(I(\mathcal{C}), \xrightarrow{ord})$ and the agreement properties.

The notation of Figure 9 is extended to include subscript 0 representing the initial state. Any instruction with a subscript of 0 does not exist. However, we can say that an instruction that does exist follows an instruction with subscript 0 in a specified order. For example, $x_0 \xrightarrow{L} y_1$ means that the instruction y_1 occurred after the initial state in the sequence L.

Lemma 3.2 *If instruction* $8[0]_i$ *is valid in* $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ *then the* ν_{c2}=read$_c(\nu_{c1})$ *operation in* $\mathcal{M}(6_i^1)$ *returns a non-zero value.*

Proof: Instructions 6_i^1, 7_i^1 and $8[0]_i$ are ordered by Local dependence order in $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$. Also by Local dependence order any other instruction i_o that contains a write operation to $r1_c$ or $p1_c$ are $i_o \xrightarrow{c} 6_{i-1}^1$ or $8[0]_i \xrightarrow{c} i_o$. $\mathcal{M}(8[0]_i)$ is necessarily 0 = read$_c(p1_c)$. This implies that $\mathcal{M}(7_i^1)$ contains write$_c(p1_c, 0)$ which implies that $\mathcal{M}(7_i^1)$ has non-zero=ν_{c3}=read$_c(r1_c)$. This implies that $\mathcal{M}(6_i^1)$ must have non-zero = ν_{c2}=read$_c(\nu_{c1})$. ∎

Lemma 3.3 *If instruction* $3[0]_i$ *is valid in* $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{p})$ *then the* ν_{p2}=read$_p(\nu_{p1})$ *operation in* $\mathcal{M}(1_i^1)$ *returns a zero.*

Proof: In $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{p})$ instructions 1_i^1, 2_i^1 and $3[0]_i$ are ordered by Local dependence order. By Local dependence order any other instruction i_o that contains a write operation to $r1_p$ or $p1_p$ are $i_o \xrightarrow{p} 1_i^1$ or $3[0]_i \xrightarrow{p} i_o$. $\mathcal{M}(3[0]_i)$ is necessarily 0 = read$_p(p1_p)$. This implies that $\mathcal{M}(2_i^1)$ contains write$_p(p1_p, 0)$ and thus that $\mathcal{M}(2_i^1)$ has 0=ν_{p3}=read$_p(r1_p)$. Thus, $\mathcal{M}(1_i^1)$ must have 0 = ν_{p2}=read$_p(\nu_{p1})$. ∎

Lemma 3.4 *At least one copy of instruction 4 is in* $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ *between instructions* 6_{i-1}^1 *and* 6_i^1, *for* $i \geq 1$.

Proof: The sequence $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ has $6_{i-1}^1 \xrightarrow{c} 10_{i-1} \xrightarrow{c} 6_i^1$ by Local dependence order. $\mathcal{M}(10_{i-1})$ contains a write$_c(\nu_{c5}, 0)$ operation where ν_{c5} contains the number that specifies the address of shared memory variable Q. Because of the initialization Q=0 prior to instruction 6_1^1 in $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$.

By lemma 3.2 $\mathcal{M}(6_i^1)$ has non-zero = ν_{c2}=read$_c(\nu_{c1})$ and thus its write$_c(r1_c, \text{nonzero})$, where ν_{c1} specifies the address of Q. Thus, some instruction that contains a write operation that writes a non-zero value to Q must follow 10_{i-1} and precede 6_1^1 in $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$. The only instruction whose meaning contains such a write operation is instruction 4. ∎

Lemma 3.5 *At least one copy of instruction 10 is in* $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{p})$ *between instructions* 1^1_{i-1} *and* 1^1_i*, for* $i > 1$*.*

Proof: Instructions $1^1_{i-1}, 4_{i_1}$ and 1^1_i are ordered by Local dependence order in $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{p})$. $\mathcal{M}(4_{i-1})$ contains a write$_p(\nu_{p4}, \text{nonzero})$ operation where ν_{p4} contains the number that specifies the address of Q.

By Lemma 3.3 $\mathcal{M}(1^1_1)$ has $0 = \nu_{p2}$=read$_p(\nu_{p1})$ and thus its write$_p(r1_p, 0)$ where ν_{p1} specifies the address of Q. Thus, some instruction that contains a write operation that writes Q=0 must follow 4_{i-1} and precede 1^1_1 in $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{p})$. The only instruction whose meaning contains such a write operation is instruction 10. ■

Lemma 3.6 *In both* $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ *and* $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{p})$ *instructions* 4_i*,* 10_i *and* 4_{i+1} *must be ordered* $4_i \xrightarrow{p} 10_i \xrightarrow{p} 4_{i+1}$ *and* $4_i \xrightarrow{c} 10_i \xrightarrow{c} 4_{i+1}$ *for any* $i > 0$*.*

Proof: By Local dependence order we have: initial state $\xrightarrow{c} 6^1_1 \xrightarrow{c} 10_1 \xrightarrow{c} 6^1_2$. By lemma 3.4 at least one instruction 4 precedes 6^1_1. By Orderable order $4_1 \xrightarrow{c} 4_2 \xrightarrow{c} 4_3 \ldots$. Thus, $4_1 \xrightarrow{c} 6^1_1 \xrightarrow{c} 10_1$. Local dependence order ensures that no copies of instruction 10 precede 6^1_1 in $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$.

By the coherence agreement property and the fact that both $\mathcal{M}(4)$ and $\mathcal{M}(10)$ contain write operations to Q, $4_1 \xrightarrow{p} 10_1$. By Local dependence order we have: initial state $\xrightarrow{p} 1^1_1 \xrightarrow{p} 4_1 \xrightarrow{p} 1^1_2$. By lemma 3.5, at least one instruction 10 is between 1^1_1. Thus, we must have $1^1_1 \xrightarrow{p} 4_1 \xrightarrow{p} 10_1 \xrightarrow{p} 1^1_2$. By Local dependence order we can extend this to $1^1_1 \xrightarrow{p} 4_1 \xrightarrow{p} 10_1 \xrightarrow{p} 1^1_2 \xrightarrow{p} 4_2$.

Continued applications of lemma 3.4, Local dependence order on $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ and coherence agreement followed by lemma 3.5, Local dependence order on $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ and coherence agreement generate the sequence: $4_1 \xrightarrow{p} 10_1 \xrightarrow{p} 4_2 \xrightarrow{p} 10_2 \ldots$ and $4_1 \xrightarrow{c} 10_1 \xrightarrow{c} 4_2 \xrightarrow{c} 10_2 \ldots$ in $(I(\mathcal{C})|p \cup I(\mathcal{C})|w, \xrightarrow{p})$ and $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$. ■

Lemma 3.7 *There is exactly one copy of instruction 4 is in* $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ *between instructions* 6^1_{i-1} *and* 6^1_i*, for* $i \geq 1$*.*

Proof: By lemma 3.4 there is at least one. Assume there exist some pair of consecutive 6^1_{i-1} and 6^1_i instructions that have two or more instruction 4's, specifically, $4_j, 4_{j+1}, \ldots$, between them in $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$, that is $6^1_{i-1} \xrightarrow{c} 4_j \xrightarrow{c} 4_{j+1} \xrightarrow{c} 6^1_i$. By Local dependence order $6^1_{i-1} \xrightarrow{prog} 6^1_i$ and because they are consecutive there are no other 6^1_k instructions between them. By the same memory part of Orderable order $6^1_{i-1} \xrightarrow{c} 10_i \xrightarrow{c} 6^1_i$ and there is no other instruction 10 between 6^1_{i-1} and 6^1_i.

By lemma 3.6 we know that $6^1_{i-1} \xrightarrow{c} 4_j \xrightarrow{c} 10_j \xrightarrow{c} 4_{j+1} \xrightarrow{c} 6^1_i$. Thus, instructions $10_j = 10_i$. By lemma 3.2 6^1_{i-1} returns a non-zero value and thus to be valid it must be preceded in $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ by an instruction 4 (called 4_* here) the only instruction that writes a non-zero value to Q, and there must be no intervening instruction 10, the only instruction that writes Q=0. Thus we have $4_* \xrightarrow{c} 6^1_{i-1} \xrightarrow{c} 4_j \xrightarrow{c} 10_j \xrightarrow{c} 4_{j+1} \xrightarrow{c} 6^1_i$. Observe that there is no instruction 10 between 4_* and 4_j and that none can be placed there. This violates lemma 3.6. ■

Lemma 3.8 *In* $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ *new items are returned (into r3) in the same order as items are produced by the producer.*

Proof: Only instructions 4 and 10 alter the contents of Q. In the ⟨produce⟩ transformation an item is produced when $\mathcal{M}(4)$ completes its write$_p(\nu_{p4}, item)$ operation, so the order of production of items is $4_1 \xrightarrow{p} 4_2 \xrightarrow{p} 4_3 \ldots$. The coherence agreement property ensures these instructions are ordered the same in $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$.

By Lemma 3.7 order $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{c})$ has $6^1_0 \xrightarrow{c} 4_1 \xrightarrow{c} 6^1_1 \xrightarrow{c} 4_2 \xrightarrow{c} 6^1_2 \xrightarrow{c} 4_3 \ldots$. By Orderable order we have $10_1 \xrightarrow{c} 10_2 \xrightarrow{c} 10_3 \ldots$. By Local dependence order we have $6^1_i \xrightarrow{c} 7^1_i \xrightarrow{c} 9_i \xrightarrow{c} 10_i$ for all $i \geq 1$ ensuring that each 10_i returns the value produced by 4_i. ■

This establishes the theorem.

■

3.4 In the absence of dependence

While the SPC system has an implementation on an Itanium platform, we show in this section that such an implementation does not exist on an Itanium-Dep platform.

First of all, we show how any transformation respecting the pseudo code in Figure 8 fails under only Itanium-Dep. Specifically, in the C-read pseudo code, the Orderable order does not always require program order between two load instructions, even if they access the same atomic variable (Q in this case). The consumption load (load of Q in 'item ← Q') may be ordered in S_c before any load instruction resulting from the spin-loop, including the load that writes a value that consequently ends the loop. This load signals to the consumer to proceed with consumption. However under Itanium-Dep, the consumer can proceed to consumption before even performing any loads or even checking the condition in the spin-loop. That is, a consumer may consume a never-produced or a previously consumed item. This 'early' load decides the value to be returned in the interpretation by the corresponding C-read. Any attempt to construct a valid total order for the producer-consumer object X will fail. The failure of this algorithm is general as we argue next.

Theorem 3.9 *There does not exist an implementation of the SPC system on an Itanium-Dep platform.*

Proof: Assume there exists an implementation β of the SPC system on an Itanium-Dep platform. Then the interpretation of any computation of the transformed SPC multiprocess will be valid for the SPC object X. By this validity requirement, the β(C-read(X)) subroutine does not return before it finds a new produced and never consumed before item. Hence, c must *wait* in β(C-read(X)) until a new produced item is available for consumption. In Itanium-Dep, the only way by which c waits for p is through the use of a spin-loop. The spin-loop has at least one load instruction, *spin-load*.

Furthermore, there has to be a load instruction after which the value to be returned by β(C-read(X)) will be decided and does not change until β(C-read(X)) returns. Call this load the *consumption load*. Obviously, the consumption load must follow in program order some of the spin-loads. We argue that Itanium-Dep can give rise to a computation that behaves as if the consumption load is completed before the spin-loop starts.

If between a spin-load and the following (in program order) consumption load there are any stores, then these stores cannot be to the same variable that the consumption load is accessing, otherwise the communicated produced item may be changed by the consumer and lost. A store that deterministically leaves the value of an atomic variable, say v, unchanged does not necessarily write the same value back to v in Itanium-Dep. Writing the same value stored in v back to v (something of the form: $v \leftarrow v$) requires a load that precedes in program order the store. What we have now in $prog\text{-}seq_{\beta(c)}$ is:

ld8 r_1=[r] ; begin spin-loads, register r has the address of v
ld8 r_2=[r]
...
ld8 r_k=[r] ; another spin-load (not necessarily the last of the loads in the loop)

ld8 r_{k+1}=[r] ; load resulting from writing v to v
st8 [r]=r_{k+1} ; writing v back to v
...
ld8 r_{k+i}=[r]; the consumption load

Itanium-Dep is insufficient to enforce any ordering between ld8 r_{k+1}=[r] and any of the k spin-loads in $(I(\mathcal{C})|c \cup I(\mathcal{C})|w, \xrightarrow{S_c})$. Hence, it is always possible for r_{k+1} to be assigned a never-produced (the item is not the result of β(P-write(X, ι))) item or an item that has already been consumed.

Finally, any store to a different variable between (in program order) the k spin-loads and the consumption load cannot restore the lost program order in S_c between the k spin-loads and the consumption load. That is, consumption can happen earlier than it should be, and we can end up with an invalid sequence of P-write and C-read operations. Hence, the transformation β cannot be an implementation, since in Figure 8, there is at least one computation in $\mathcal{C}'$ that cannot be in $\mathcal{C}$. ■

References

[1] A. Adir, H. Attiya, and G. Shurek. Information-flow models for shared memory with an application to the PowerPC architecture. *IEEE Trans. on Parallel and Distributed Systems*, 14(5):502–515, 2003.

[2] H. Attiya and R. Friedman. Programming DEC-Alpha based multiprocessors the easy way. In *Proc. 6th Int'l Symp. on Parallel Algorithms and Architectures*, pages 157–166, 1994. Technical Report LPCR 9411, Computer Science Department, Technion.

[3] P. Chatterjee and G. Gopalakrishnan. Towards a formal model of shared memory consistency for intel itaniumtm. In *Proc. 2001 IEEE International Conference on Computer Design (ICCD)*, pages 515–518, Sept 2001.

[4] M. Herlihy and J. Wing. Linearizability: A correctness condition for concurrent objects. *ACM Trans. on Programming Languages and Systems*, 12(3):463–492, July 1990.

[5] L. Higham, L. Jackson, and J. Kawash. Specifying memory consistency of write buffer multiprocessors. *ACM Trans. on Computer Systems*. In Press.

[6] L. Higham, L. Jackson, and J. Kawash. Programmer-centric conditions for Itanium memory consistency. Technical Report 2006-838-31, Department of Computer Science, The University of Calgary, July 2006.

[7] L. Higham and J. Kawash. Tight bounds for critical sections on processor consistent platforms. *IEEE Trans. on Parallel and Distributed Systems*. In Press.

[8] L. Higham and J. Kawash. Critical sections and producer/consumer queues in weak memory systems. In *Proc. 1997 Int'l Symp. on Parallel Architectures, Algorithms, and Networks*, pages 56–63, December 1997.

[9] L. Higham and J. Kawash. Java: Memory consistency and process coordination (extended abstract). In *Proc. 12th Int'l Symp. on Distributed Computing, Lecture Notes in Computer Science volume 1499*, pages 201–215, September 1998.

[10] L. Higham and J. Kawash. Memory consistency and process coordination for SPARC multiprocessors. In *Proc. of the 7th Int'l Conf. on High Performance Computing, Lecture Notes in Computer Science volume 1970*, pages 355–366, December 2000.

[11] Intel Corporation. A formal specification of the intel itanium processor family memory ordering. http://www.intel.com/, Oct 2002.

[12] Intel Corporation. *Intel Architecture Architecture Software Developer's Manual: Volume 3: Instruction Set Reference*. Intel Corporation, Oct. 2002.

[13] Intel Corporation. Intel itanium architecture software developer's manual, volume 2: System architecture. http://www.intel.com/, Oct 2002.

[14] L. Jackson. Complete framework for memory consistency with applications to Itanium multiprocessors. Ph.D. Dissertation, in preperation.

[15] R. Joshi, L. Lamport, J. Matthews, S. Tasiran, M. Tuttle, and Y. Yu. Checking cache-coherence protocols with tla, 2003.

[16] L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. on Computers*, C-28(9):690–691, September 1979.

[17] Y. Yang, G. Gopalakrishnan, G. Lindstrom, and K. Slind. Analyzing the Intel Itanium memory ordering rules using logic programming and SAT. In *Proc. 12th Advanced Research Working Conference on Correct Hardware Design and Verification Methods*, 2003.