

Exploring the Performance Limits of Out-of-order Commit

Mehdi Alipour, Trevor E. Carlson and Stefanos Kaxiras
Uppsala University, Department of Information Technology, Uppsala, Sweden
FirstName.LastName@it.uu.se

ABSTRACT

Out-of-order execution is essential for high performance, general-purpose computation, as it can find and execute useful work instead of stalling. However, it is limited by the requirement of visibly sequential, atomic instruction execution — in other words *in-order instruction commit*. While in-order commit has its advantages, such as providing precise interrupts and avoiding complications with the memory consistency model, it requires the core to hold on to resources (reorder buffer entries, load/store queue entries, registers) until they are released in program order. In contrast, *out-of-order commit* releases resources much earlier, yielding improved performance with fewer traditional hardware resources. However, out-of-order commit is limited in terms of correctness by the conditions described in the work of Bell and Lipasti. In this paper we revisit out-of-order commit from a different perspective, not by proposing another hardware technique, but by examining these conditions one by one and in combination with respect to their potential performance benefit for both non-speculative and speculative out-of-order commit. While correctly handling recovery for all out-of-order commit conditions currently requires complex tracking and expensive checkpointing, this work aims to demonstrate the potential for selective, speculative out-of-order commit using an *oracle implementation without speculative rollback costs*. We learn that: a) there is significant untapped potential for aggressive variants of out-of-order commit; b) it is important to optimize the commit depth, or the search distance for out-of-order commit, for a balanced design: smaller cores can benefit from shorter depths while larger cores continue to benefit from aggressive parameters; c) the focus on a subset of out-of-order commit conditions could lead to efficient implementations; d) the benefits for out-of-order commit increase with higher memory latency and works well in conjunction with prefetching to continue to improve performance.

ACM Reference format:

Mehdi Alipour, Trevor E. Carlson and Stefanos Kaxiras. 2017. Exploring the Performance Limits of Out-of-order Commit. In *Proceedings of CF'17, Siena, Italy, May 15-17, 2017*, 10 pages.

DOI: <http://dx.doi.org/10.1145/3075564.3075581>

1 INTRODUCTION

Dynamically-scheduled superscalar processors execute instructions out-of-order but *commit in-order* to present to the programmer the

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CF'17, Siena, Italy

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. 978-1-4503-4487-6/17/05...\$15.00

DOI: <http://dx.doi.org/10.1145/3075564.3075581>




Figure 1: Performance comparison (harmonic mean of IPCs across SPEC CPU2006) of in-order and two types of safe out-of-order commit, reluctant (ROOC) and aggressive (AOOC), with a commit depth of 4 and ROB size for increasingly aggressive microarchitectures. These experiments respect all traditional commit conditions, and show that aggressive out-of-order commit can reach the performance of the next class of processor.

illusion that instructions execute *atomically* and *sequentially* as intended by the program. In this context, precise interrupts are easily provided as the processor verifies correct execution before each instruction is committed [15].

The disadvantage of in-order commit is that it ties up resources (reorder buffer [ROB] entries, load/store queue [LSQ] entries, and physical registers) for a much longer time than is necessary for correct execution. In-order commit ties up resources until all instructions complete and commit in the correct sequential program order. This means that execution is halted when any of the resources are exhausted: when the ROB fills up, or when we run out of either LSQ entries or registers. To overcome this hurdle, designers size these structures so that they minimize the chances of any single resource becoming exhausted, creating a *balanced* microarchitecture. This, however, contributes to the power-inefficiency of the OoO cores. At the very least, after a certain point, increasing the size of the LSQ (which is usually implemented as an expensive CAM) or the size of the register file results in an increase in energy consumption that far exceeds the performance benefit, a significant disadvantage for power-constrained processors.

Thus, the incentive for pursuing *out-of-order commit* (OOC) lies in the promise of higher performance with fewer resources. A turning point in our understanding of out-of-order commit came with the work of Bell and Lipasti [3] who articulated the limiting factors for OOC. The necessary conditions to allow an instruction to be committed range from the completion status of the instruction itself, to the branch prediction and exception state of intervening instructions. Several proposals for out-of-order commit, implicitly or explicitly, abide by these conditions, providing the means *and resulting efficiency cost* to enforce them.

What we do do: The question we explore in this paper is what could the performance gain be if we had the means to evade any one or even all of these conditions together. This work is the first, to our understanding, to evaluate the potential performance contribution of each individual commit condition. For example, what would we gain if we could correctly commit past unresolved branches, or stores with unresolved addresses, or instructions that could generate an exception? Answering this question allows us to focus on understanding the most profitable conditions.

What we do not do: However, in this paper we do not quantify the *cost* that would be required to guarantee *correctness* when committing past any or all of these conditions, as this would be tied to a specific hardware or software mechanism. Instead, we evaluate the potential performance benefits available as the means to gauge the efficiency of future proposals in relation to their cost. In the same vein, we have not implemented every possible core optimization that could be construed as extending the instruction window outside the core (e.g., run-ahead execution) as we would need to take into account the additional cost in all the possible ways that it could be reallocated (e.g., increasing cache size or improving the branch predictor). Clearly this is intractable. Thus, we chose to study three *known* architectures to allow for *known* baselines with *understood* costs. Adding proposed processor improvements (and evaluating their power, energy, area, etc.) is beyond the scope of this work.

Contributions: We study *safe* out-of-order commit (refraining from committing until all conditions are met) and *unsafe* out-of-order commit (committing even before all conditions are met assuming the potential for correct recovery and commit). In addition to studying both safe and unsafe out-of-order commit as minimum and maximum potential performance improvement, we also study out-of-order commit in two additional dimensions: aggressiveness and degree of speculation:

- *Aggressiveness:* A related dimension that determines this potential benefit is the aggressiveness of the out-of-order commit implementation. Prior work [14] examines this aggressiveness in their implementation of checkpoint-based out-of-order commit, where they show a middle-ground in aggressiveness is needed to mitigate checkpoint-based penalties while still improving performance. In addition, work on non-speculative out-of-order commit [3] has recognized that a restricted commit window can achieve a good portion of the performance improvements compared to an unrestricted, unlimited out-of-order commit implementation.
- *Degree of speculation:* Apart from the aggressiveness of commit, we also consider varying degrees of speculation. The insight of this study is that selective speculation, preserving or relaxing of one more conditions necessary for out-of-order commit, could lead to more efficient implementations. In fact, recent work [9, 10] has shown that this might be the case for load instructions.

Results: This study aims to provide a guide post for each major parameter of out-of-order commit to provide an upper bound on the performance benefits given aggressiveness and degree of selective speculation. In this work, we show that

- there is significant untapped potential for *unsafe* out-of-order commit beyond traditional in-order and *safe* out-of-order commit; the gap widens in more powerful architectures;
- for energy-efficient cores with moderately-sized instruction windows, reluctant, limited, out-of-order commit is sufficient to reap the most benefit, while aggressive versions of out-of-order commit become a requirement for larger cores;
- the order of importance of the commit conditions changes depending on the type of application (e.g., FP vs. integer), the architecture (limited or aggressive OoO processor), and the aggressiveness of out-of-order commit;
- a focus on specific out-of-order commit conditions could be an important future direction for high-performance, efficient out-of-order processors;
- potential benefits of out-of-order commit increase with memory latency (relatively more for unsafe) while prefetching benefits are orthogonal to out-of-order commit benefits. This raises the enticing possibility of reducing system-wide financial cost without compromising performance by coupling dense but higher-latency (slow-but-efficient) DRAM with out-of-order commit cores.

All in all, our study shows that as a future direction, unsafe out-of-order commit appears to be very promising. This is especially true if the potential benefits can be tapped with much more efficient and selective mechanisms to guarantee correctness, instead of bulk checkpointing and rollback. The rest of this paper is organized as follows. In **Section 2** we first provide an overview of the conditions that need to be honored for in-order commit as well as provide an overview of out-of-order commit. Next, in **Section 3** we present our evaluation methodology and simulated system configurations. **Section 4** provides a detailed performance analysis for each out-of-order commit condition, both from the point of view of aggressive and reluctant out-of-order commit. Finally, the related works and conclusion are presented in **Section 5** and **6** respectively.

2 OUT-OF-ORDER COMMIT

The introduction of the reorder-buffer (ROB) to provide in-order commit in an out-of-order scheduled superscalar pipeline was an important advance for computer architecture culminating an effort that started with Tomasulo’s algorithm, and included techniques such as reservation stations and the register-update unit (RUU) [16]. The reorder buffer maintains precise architectural state in the presence of interrupts, unknown memory dependencies, or memory re-orderings that can perturb the ordering required by a memory consistency model.

Out-of-order commit, on the other hand, attempts to break this rigid updating of the architectural state either in a *safe* way (i.e., one that does not require *additional* speculation and rollback to revert changes to the architectural state) or in an *unsafe* way (i.e., one that does).

Safe vs. Unsafe OOC. A turning point in our understanding of out-of-order commit came with the work of Bell and Lipasti [3] in the form of a number of limiting conditions for *safe* out-of-order commit. The necessary conditions to allow an instruction to be committed out-of-order are:

- (1) The instruction is completed: instructions can commit only after their completion.
- (2) The instruction is not involved in memory replay traps. This condition simply says that we cannot commit speculative loads or their dependent instructions. This condition relates to unresolved memory dependencies or memory consistency enforcement. For example, total store order (TSO) requires a replay of speculative loads that violate load→load ordering when this reordering is detected by other cores.
- (3) Register WAR hazards are resolved (i.e., a write to a particular register cannot be permitted to commit before all prior reads of that architectural register have been completed).
- (4) Previous branches are successfully predicted. This condition simply says that we can commit only while on the correct path of execution.
- (5) No prior instruction in program order is going to raise an exception. This condition provides precise interrupts and is essential in easing handling of, e.g., page faults.

In the rest of this paper we will use the following convention to discuss how we *evade* the Bell/Lipasti conditions.

- (1) **Safe.OOC** where all out-of-order-commit conditions are preserved. This case provides the minimum potential performance improvement of out-of-order commit, but also the minimum hardware to implement as it does not rely on speculation and rollback beyond what is already available in the out-of-order core.
- (2) **Unsafe.OOC** where one or more (or all) of the out-of-order commit conditions are evaded (apart from true dependencies). Doing so, the maximum potential performance improvement of out-of-order commit is evaluated, but this may require extra support for speculation and rollback to be able to revert changes in the architectural state that were found to be incorrect after the commit.

Reluctant vs. Aggressive OOC Aside from the limiting conditions described above, a separate dimension is the aggressiveness of committing out-of-order.

Thus, concerning the mechanics of out-of-order commit, we distinguish two versions:

- (1) **Reluctant** out-of-order-commit (ROOC): where the out-of-order commit mechanisms are engaged only *when needed* and,
- (2) **Aggressive** out-of-order-commit (AOOC): where the out-of-order commit mechanisms are continuously active, looking for opportunities to commit instructions as early as possible.

While the *always on* nature of the aggressive out-of-order commit is obvious in its meaning, the *when needed* of the reluctant out-of-order-commit requires clarification. For this paper, reluctant out-of-order commit is engaged only when the core is in imminent danger of going into a complete stall. In other words, we engage reluctant out-of-order commit only when one of the critical resources (ROB entries, registers, load-store queue entries) is all but exhausted and cannot support the fetching of new instructions in the front end of the pipeline. As such, reluctant out-of-order commit acts as a safety valve to release the pressure on resources (just before this

pressure reaches a critical point), rather than aggressively trying to keep resource pressure low.

In contrast, aggressive out-of-order commit releases resources more eagerly, but disregards the following issues:

- (1) it might prove wasteful as traditional in-order commit may still be able to provide sufficient resources for forward progress;
- (2) it may be futile as the chances of encountering an instruction that restricts further commit (e.g., an unresolved branch) tends to increase with aggressiveness.
- (3) it creates a significant management problem as out-of-order commit can create gaps in several structures, including the ROB and also the load queue and store queue (which is not completely addressed in prior works [3, 13, 14]).

Commit Width and Depth

The final parameter to explore within the context of out-of-order commit is the commit depth to scan for potential instructions to commit out-of-order. While commit width is the number of instructions that can be committed simultaneously per cycle, the commit depth is the measure of how far the core can scan looking for instructions to commit out-of-order in a given cycle.

2.1 OOC Conditions

In Section 3, we describe the methodology used to commit instructions out-of-order. With this methodology, we can selectively relax the commit conditions described in Section 1 (except completion) while still guaranteeing correct execution. For example, by relaxing the branch condition (Unsafe_BR): “committing only down a non-speculative path,” we can continue to free resources past unresolved branches but effectively only commit from the correct path. In this paper we do not evaluate the implementation required to relax these conditions, but instead evaluate the potential for performance improvement. We evaluate the maximum potential performance improvement with a speculative out-of-order rollback cost of zero. In this section we provide details on the performance implications of the relaxation of a single and combination of conditions.

2.1.1 Instruction complete. The core waits for an instruction to finish executing before commit can occur. We do not examine early commit of loads [9, 10] that miss in the cache and instead we consider them available for commit only after the data returns and is bound to the destination register.

2.1.2 Memory replay traps (safe_ST and safe_LD). We describe two sub-cases for this condition:

Store-Load (safe_ST): This condition applies to same-thread memory dependencies involving a store and a load. In particular, we cannot commit a load out-of-order in the presence of a prior store with an unresolved address. If the store and the load prove to be dependent (the load should have taken the value of the store) the commit would have been incorrect. The LD condition disallows the commit of a load and its dependent instructions until all prior stores resolve their addresses and all the memory dependencies are correctly enforced. By relaxing this condition, we can commit loads and their dependent instructions even if prior non-aliasing stores have unresolved addresses.

Load-Load (safe_LD): This concerns memory consistency models that enforce load→load ordering (e.g., Sequential Consistency

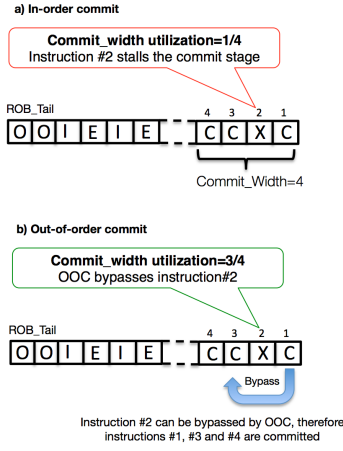


Figure 2: Conceptual comparison of in-order and out-of-order commit (commit depth=4). C: Ready to commit, X: Not ready to commit, E: Executing, I: Issued, O: Empty entry.

or TSO). Under this ordering constraint it is possible to allow loads out-of-order as long as this is *not observed in the memory system*. The safe_LD condition disallows the out-of-order commit of loads unless it is guaranteed that the correct order will be observed by the memory system. To relax this condition we allow load→load re-orderings that are not observed by other cores. A very specific case would be a memory mapped IO (MMIO) request that might change the order of memory operations. The MMIO case acts as a ‘coprocessor’, meaning that we have a multi-processor system here. We ignore memory requests from other cores (IO/Coprocessor).

2.1.3 WAR hazards. WAR hazards are already handled by the out-of-order core within the ROB, and we assume a solution such as the Value Buffer [13] for committing out-of-order. Thus, we do not consider this condition further.

2.1.4 Unresolved Branches (safe_BR). This condition guarantees that we commit only from the correct path of execution. Out-of-order commit should not proceed past unresolved branches until they are correctly resolved. We can relax this condition and commit past an unresolved branch if we are able to undo the commit. To evaluate maximum performance potential we assume a zero *roll-back* cost for out-of-order commit misspredictions. However, the normal branch misprediction cost (10 cycles) is faithfully accounted (see Subsection 4.9 for more details.). In addition, we evaluate the *rollback count* for this condition in Subsection 4.9.

2.1.5 Exceptions (safe_EXC). This condition caters to precise interrupts. Enforcing this condition requires that we do not commit past an instruction (floating-point, memory access, or any instruction that may cause an exception) unless we make sure that the instruction will not cause an exception. To relax this safe_EXC condition, we assume the code regions are exception free.

2.2 Safe and Unsafe Out-of-Order Commit

Normally, discussion of out-of-order commit tends to focus on the changes that occur in the back end of the pipeline. However, the purpose of out-of-order commit is to enable the *front end* to proceed. The perception that out-of-order commit increases performance can also be misleading: *instruction execution is not sped up*; it is

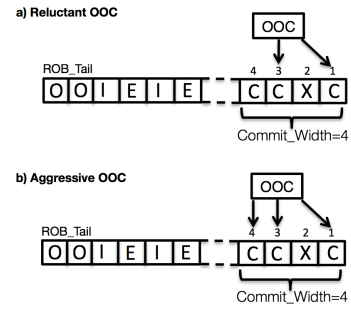


Figure 3: Functionality of AOOC and ROOC. In this example AOOC commits one more instruction than ROOC (commit depth=4). C: Ready to commit, X: Not ready to commit, E: Executing, I: Issued, O: Empty entry.

the removal of conditions that stall the front end that increases performance. More specifically, in an *unconstrained* architecture (unrestricted ROB, registers, and load/store queue entries), *safe* out-of-order commit does not perform faster than in-order commit. For restricted, real-world implementations, Safe_OOC helps to ameliorate this problem.

In contrast, Unsafe_OOC has the capacity to exceed the performance of an unconstrained in-order commit architecture as it removes penalties needed to guarantee correctness (e.g., correctly handling memory dependencies, correctly enforcing memory consistency ordering, etc.) in situations that the hardware does not have any other means of imposing such correctness. In this case, Unsafe_OOC has the potential to violate correctness and requires a means to revert back to a safe state if a violation occurs. This can be a good trade-off when the conditions that violate correctness are rare. We implement an oracle version of Unsafe_OOC, in that it violates the correctness conditions because it will be safe to do so (and therefore no correctness problems will occur). This removes the need to provide the mechanisms to revert to a safe state in the case of a misspeculation. This model provides a best-case speedup as recovery from misspeculation is *zero cost*.

2.3 Aggressive and Reluctant OOC

Orthogonal to the enforcement of the out-of-order commit conditions, the aggressiveness of out-of-order commit plays a significant role in the resulting performance and the cost it incurs. We introduced two approaches for out-of-order commit: Aggressive (AOOC) and Reluctant (ROOC) out-of-order commit. A good way to describe them is to contrast their main difference: *how often each mechanism is engaged*.

AOOC is engaged all the time, i.e., it constantly tries to find instructions to commit out-of-order if the opportunity arises. In this respect it aims to increase commit bandwidth.

In contrast, ROOC is only engaged when one of the critical resources in the core (ROB entries, registers, Load/store queue entries) is about to be exhausted. ROOC is concerned about front end stalls, not about commit bandwidth. This means that the number of instructions that ROOC needs to find that can commit out-of-order is *limited*: ROOC needs to provide enough free entries in the ROB/registers/LSQ so that the front end can dispatch as many instructions as possible (up to the dispatch width) to the out-of-order engine. Figure 2 shows this with an example. The reason that the

Table 1: Baseline core parameters

Full system simulation	3.4 GHz, x86-64
L1i/d	32KiB, 8-way, 4clk
L2	256KiB, 8-way, 12clk
L3	1MiB, 8-way, 36clk
DRAM	200clk
Branch Predictor	Tournament, front end penalty 10clk
Prefetcher OFF(default)/ON	Stride, degree 8

Table 2: Microarchitecture configuration with reorder buffer (ROB), instruction queue (IQ), load and store queues (LQ/SQ) and integer and floating-point register file (RF) details. Dispatch width (D), commit width (CW) and Out-of-Order commit depth (CD) is the same to enable a fair comparison (SLM hardware has a D/CW of 2). Register values are additional physical registers above architectural state.

Microarchitecture	D/CW/CD	ROB	IQ	LQ/SQ	RF(INT,FP)
Silvermont-like (SLM)	4/4/8	32	32	10/16	32,32
Nehalem-like (NHM)	4/4/8	128	56	48/36	68,68
Haswell-like (HSW)	4/4/8	192	60	72/42	130,130

commit stage of an in-order commit is blocked is because of an unresolved instruction at the head of ROB. A four-wide superscalar, shown in Figure 2, tries to commit four instructions per cycle. While the first instruction at the head of ROB can commit, the second oldest instruction causes the commit stage to block and instead of four, only one instruction commits (in-order). The behavior of out-of-order commit depends on its aggressiveness:

AOOC: AOOC attempts to find up to *commit-width* instructions so it can satisfy the need to commit at the highest possible commit bandwidth. In our example it aims to find four instructions. However, since only one more is needed to fill the gap in the head group of four leading instructions, AOOC produces an excess of commit-ready instructions that can potentially be exploited in the near future.

ROOC: ROOC on the other hand, aims to find the minimum number of instructions needed to commit (out-of-order) so that no resource is exhausted and the front end can continue to issue instructions at its peak bandwidth. The reason for seeking the *minimum* number of instructions to commit out-of-order is that this minimizes the perturbations in instruction order. This could potentially lead to more efficient hardware implementations. While Figure 3 only considers the ROB, in the general case all dispatch-related resources are considered in the same way.

In our particular example, ROOC needs to find just one instruction to commit out-of-order. The ROB already contains two empty slots, and the in-order commit mechanisms can find one instruction to commit, leaving one left for ROOC. In contrast, in AOOC mode, there are in total three instructions that commit (instructions #1, #3 and #4 in Figure 3). The result is that AOOC needs to scan deeper than ROOC.

3 METHODOLOGY

We use the gem5 [4] simulator in *full system* mode to simulate an x86-64 target with a frequency of 3.4GHz. To test our models, we use the SPEC CPU2006 benchmark suite. We use ten uniformly distributed checkpoints from each benchmark. Each simulation checkpoint has three phases that begin with 250M instructions of cache warm-up, followed by 100k instructions of detailed pipeline




Figure 4: Commit bandwidth distribution for the SPEC CPU2006 benchmarks of a 4-wide core with in-order commit and out-of-order commit respecting all commit conditions (Safe.OOC). Out-of-order commit increases commit pressure even without aggressive speculation.

warm-up and ends with a detailed simulation of 100M instructions. Furthermore, three different configurations, similar to three conventional out-of-order processors, are used: Intel’s SLM, NHM, and HSW [5] microarchitectures. Tables 1 and 2 list the detailed configuration of the simulation environment.

To implement our out-of-order commit model, we first configure a simulated machine with a very large number of core resources. We then monitor the number of committed and non-committed instructions that appear in the pipeline, and control, at dispatch, whether we can support additional instructions in the back end of the processor. In this way, we dynamically determine the resource availability in the processor for each cycle based on the out-of-order commit conditions.

4 OUT-OF-ORDER COMMIT EVALUATION

In this section we analyze the benefits of out-of-order commit on the performance of a number of applications. We look at how the commit bandwidth changes with out-of-order commit, and how the effective resource size of each critical component changes as we enable different out-of-order commit conditions. Next, we show how out-of-order commit conditions affect performance. Safe.OOC and Unsafe.OOC are two extreme points that are defined by either enabling (respecting) or disabling all of the out-of-order commit conditions. This results in a minimum and maximum potential performance improvement across the benchmark suite. To understand the effect of each condition in isolation, we study the effect of each one, both in the presence and absence of other conditions. These studies on Safe.OOC and Unsafe.OOC conditions were conducted for both Aggressive (AOOC) and Reluctant (ROOC) out-of-order commit.

4.1 Microarchitecture Aggressiveness

We target three microarchitectures resembling Intel’s Silvermont (SLM), Nehalem (NHM) and Haswell (HSW) as small, medium and large cores (See Table 2 for details). As an overview, Figure 1 shows the performance improvement for each microarchitecture assuming Safe.OOC (all conditions respected) for all benchmarks on average. We can see that in the case of narrow commit depth (four in this figure), a relatively small out-of-order processor (SLM), has more potential for relative improvement compared to the medium and aggressive microarchitectures. The reason is that the smaller




Figure 5: Effective resource use in in-order and Aggressive out-of-order commit across SPEC CPU2006.

processor (with a shorter instruction reach and, given a balanced design, smaller hardware structures) will more likely stall as it exposes a smaller amount of the potential ILP in an application. In case of a larger commit depth, the more aggressive cores (NHM and HSW) have higher potential performance improvement (See Section 4.4 for a detailed overview). Out-of-order commit frees the processor from the traditional limits, reducing the number of times the processor experiences exhausted resources. In medium and large aggressive cores, thanks to a larger ROB as well as other hardware resources, more intrinsic ILP is extracted by traditional in-order commit, leaving less potential for out-of-order commit with a narrow commit depth.

4.2 The Effect on Commit Bandwidth

Figure 4 shows the distribution of the number of committed instructions per cycle for three different microarchitectures, for both in-order and out-of-order commit, across all SPEC CPU2006 benchmarks. Although an in-order commit, 4-wide commit HSW microarchitecture can retire up to four instructions per cycle, this occurs, on average, less than 20% of the time. In practice we see a large number of commit stage stalls (zero instructions committed per cycle). For out-of-order commit, the distribution shifts toward four instructions per cycle reflecting the improved commit performance (and the resulting improvement in overall performance). Finally this figure shows that for smaller, less aggressive cores (such as SLM, see Table 2), out-of-order commit provides a relatively larger improvement compared to the other microarchitectures because of the short commit depth of four. Larger cores improve more with more aggressive commit depths and out-of-order commit parameters (see section 4.4 for details).

4.3 The Effect on Resources

One of the main issues that out-of-order commit can help resolve is the early release (and subsequent reuse) of hardware resources that otherwise would still be required to maintain the in-order state in an in-order commit processor. Therefore, with a small ROB (and other appropriately sized structures) given in-order commit, the core can more easily stall when it runs out of resources. For example, consider a ROB size of 32 from a SLM microarchitecture. When the ROB is full it contains 32 micro-operations in flight and the CPU’s back end will no longer accept additional micro-operations from the front end. Increasing the physical size of the ROB along with other resources is one potential, but rather expensive, solution to this issue. An alternative solution, and one of the benefits of using

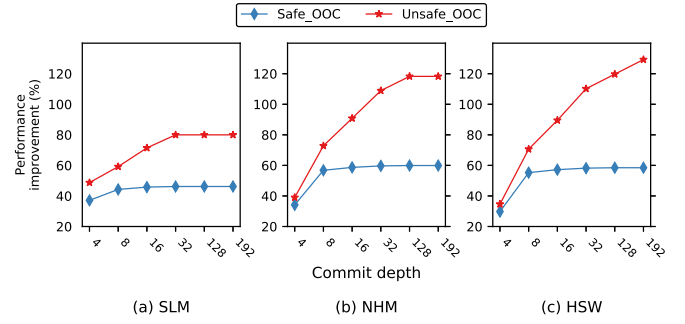


Figure 6: Effect of commit depth on performance improvement normalized to their respective in-order commit performance across SPEC CPU2006. The smaller core receives less benefit from increasing the depth of commit. Safe_OOC saturates at a commit depth of 8, 16 and 32 for SLM, NHM and HSW respectively. (the distance to the first unresolved instruction from the head of the ROB). There is no saturation in performance improvement of Unsafe_OOC while the commit depth is increased.

out-of-order commit, is the early release of resources. This early release increases the *effective* size of a resource (compared to an in-order commit core) improving performance of the core.

Benchmark specific analysis. The xalan benchmark is one of the top 5 benchmarks with a large number of CPU stalls caused by exhausted resources. Both AOOC and ROOC are very effective for the xalan benchmark. OOC is able to provide additional free entries in the ROB, RF and LSQ (see Figure 7). On the other hand, les1ie3d has the lowest number of CPU stalls based on exhausted resources, which limits the potential for improvement with OOC.

Figure 5 compares the effective size of the SLM microarchitecture between in-order commit and both safe and unsafe aggressive out-of-order commit (AOOC) models for the SPEC CPU2006 benchmarks. The results are normalized to a fixed size SLM_IOC microarchitecture (see Table 2). An effective size of 1.0 translates to frequent stalls due to resource exhaustion, while sizes greater than 1.0 shows the effective resource size increases due to OOC. For aggressive out-of-order commit, the larger effective sizes show the reach of this technique. We see that the utilization of all structures except for the ROB is almost the same for safe and unsafe out-of-order commit, which allows Safe_OOC to achieve most of the performance of the unsafe version. Nevertheless, the unsafe core is much better at improving the reach of the ROB, allowing applications like hmmer continue to show a benefit when moving from safe to unsafe out-of-order commit. See Section 4.5.2 for more details on hmmer.

4.4 Evaluation of Commit Depth

To gauge the effect of the commit depth (i.e., how far we scan the ROB to find instructions that can commit out-of-order) we impose a hard limit on it and evaluate the effects on the resulting performance. The strictest limit is the commit-width itself: starting to commit out-of-order from the first *commit-width* instructions. We then relax this to the immediate vicinity (e.g., double the commit-width) and progressively relax until we reach the size of the ROB.

In Figure 6 we see that a commit depth of 4 (equal to commit width) provides the smallest benefit, but also the smallest difference

between Safe_OOC and Unsafe_OOC. In addition, the SLM core benefits more from OOC compared to NHM and HSW when commit depth is smaller than 8. At a commit depth of 8 and above, the large aggressive cores benefit much more than the smaller core type with the maximum improvement of Unsafe_OOC at 80%, 119% and 129% for SLM, NHM and HSW, respectively. For aggressive cores, the larger commit depth allows for continued performance improvements, and will be necessary for the design of a balanced, aggressive out-of-order commit design.

4.5 Out-of-Order Commit Performance

In this section we analyze the out-of-order commit conditions to determine the minimum (and maximum) performance improvement potential. The minimum and maximum improvement is provided by Safe_OOC and Unsafe_OOC, respectively. In Figure 7, we show the amount of improvement provided by both safe and unsafe out-of-order commit, for both aggressive and reluctant modes, for all three microarchitectures.

4.5.1 Safe_OOC. Honoring all out-of-order commit conditions results in a modest performance improvement. One implication for future processors is that Safe_OOC does not require additional support for speculative out-of-order rollback recovery mechanisms; it requires support for the commit of instructions out-of-order and the freeing of structures for use by future instructions.

Safe_AOOC. When Safe_AOOC is evaluated on the SLM microarchitecture (see Figure 7a), the range of improvement spans from a low of 3% (leslie3d) up to 82% (xalan), with an average of 44%. In the NHM microarchitecture, the improvement ranges from 9% to 108% for milc and xalan, with an average improvement of 57%, and for HSW, we see an improvement of 1% for mcf to 110% for wrf, with an average of 55%. (See Section 4.4 for more details).

Safe_ROOC. For Safe_ROOC, the performance improvement is lower for all three microarchitectures compared to AOOC (See Subsection 2.3 for additional details). In the SLM microarchitecture, the range of performance improvement of safe ROOC is from 2% to 55% for calculix and gcc respectively, with an average improvement of 20%. In the NHM microarchitecture, we see performance improvements that range from 1% (mcf) to 22% (bwaves) with an average improvement of 7% (8% for HSW). Because NHM and HSM have fewer CPU stalls, ROOC has less of an effect on performance compared with the smaller, more efficient core.

4.5.2 Unsafe_OOC. By relaxing all conditions, Unsafe_OOC provides the maximum potential for performance improvement. Unsafe_OOC will require recovery mechanisms for these techniques, which can reduce the performance potential because of recovery costs. To understand the effectiveness of all conditions together we consider zero cost for recovery for any misspeculated out-of-order commit condition.

Unsafe_AOOC This technique provides the highest performance improvement for all three different architectures. In SLM cores, the improvement ranges from 9% to 120% for libquantum and astar applications respectively and the average is 59%. In NHM architecture, the average is 72% and the range is between milc and astar respectively with 11% and 196% improvement. In HSW cores, similar to NHM cores, astar has the maximum benefit from Unsafe_ooc with 192% improvement, while the minimum improvement is for

dealii, at 7%. The average improvement for this class of architecture is 70%.

Unsafe_ROOC Reluctant out-of-order commit is lower performing because it is not continuously looking to commit additional instructions (See Section 2.3 for details). In the case of SLM, the improvement ranges from 3% to 93% respectively for leslie3d and xalan with an average of a 38% improvement. In NHM, mcf and bwaves with 1% and 22% show the maximum and the minimum improvement with an average of 7% (HSW is similar with an average of 8%). Between the three microarchitectures, the limited SLM benefits the most from ROOC because of the large number of stalls seen by this core. Therefore ROOC, especially Unsafe_ROOC, is an interesting methodology to improve the performance of relatively small but energy efficient CPUs as we see a relatively high performance improvement for a less aggressive commit implementation.

Benchmark-specific Analysis. The hammer benchmark is a particularly strong case for the benefits of out-of-order commit for SLM. This application is L1-cache resident, and exhibits very few last-level cache misses. Nevertheless, we still see a very strong improvement in performance, from 33% to 52% for Safe_OOC, increasing to 51% to 66% for Unsafe_OOC. Looking ahead to Figure 8 and Figure 11, we can see that evading the branch condition provides the most benefit for this application. Making room for additional instructions to allow the hardware to expose additional ILP works well even for those applications without a significant number of LLC misses.

The mcf benchmark contains load-dependent branches, has the highest MPKI (misses per kilo instructions) among the benchmarks and therefore, it has a rather low IPC when it is executed on an in-order commit CPU. This results in a good opportunity to improve performance as it is extremely limited by these misses.

4.6 Performance Effects of Commit Conditions

In the previous section, by analyzing safe and unsafe out-of-order commit, we observe that there is a large gap between the performance improvement of these two implementations. Understanding the cause of this performance improvement (by looking at individual commit conditions in isolation), allows us to better understand where to focus future hardware efforts.

4.6.1 Positive Contribution of Out-of-Order Commit Conditions. To study the gap between safe and unsafe out-of-order commit (Figure 7), we analyze the effect of relaxing each condition in the presence of the other preserved conditions in Figure 8. We analyze the SLM microarchitecture in detail and provide averages across all microarchitectures for both AOOC and ROOC. Each out-of-order commit condition is analyzed in isolation, and we consider Unsafe_OOC (all relaxed conditions) as the 100% *potential improvement budget*. In the case of the mcf benchmark in Figure 7a, the safe and unsafe OOC performance improvement is 33% and 71% respectively (46% of the potential improvement budget is provided by Safe_OOC). We also observe that by relaxing the LD condition (unsafe_LD), 52% of potential improvement budget is achievable (see Figure 8a). In Figure 8, we can see in some applications (like namd in AOOC mode and leslie3d in ROOC mode) that relaxing just a single condition is not sufficient to fill the gap between safe and unsafe OOC. This does not mean that a single condition




Figure 9: A comparison between safe and unsafe out-of-order commit across DRAM latencies for SLM and AOOC on SPEC CPU2006. OOC improvement increases with a higher DRAM latency. Unsafe_OOC outperforms Safe_OOC. This study uses caches that are four times smaller compared to Table 1 to put additional stress on the DRAM. This has been done for both in-order and out-of-order commit configurations.




Figure 10: The effect of aggressive out-of-order commit on the performance with and without prefetchers in the L1 data cache across SPEC CPU2006. Improvement is relative to the baseline IOC architectures without prefetchers.

4.7 Memory Latency Evaluation

Current DRAM cells are optimized for cost, and not for access latency [12]. The potential to use more affordable, but higher latency memory could have a large impact on allocation of memory in data-center servers as high density DRAM modules cost much more than those that are $4\times$ smaller ($1.75\times$ per GB [2]). To evaluate the potential of out-of-order commit to handle higher DRAM latency, we evaluate memory latency from 200 cycles to 1000 cycles in 200 cycle increments. In addition, we reduce the size of all caches by 4 times when compared to Table 1, to put additional stress on the DRAM subsystem. Here we see an increasing performance improvement for the Unsafe_OOC condition, while Safe_OOC increases linearly compared to the in-order commit core. For memory-intensive applications, an efficient implementation of Unsafe_OOC could allow the use of denser, higher-latency DRAM that could potentially cost much less.

4.8 Prefetching Evaluation

Prefetching is an important way to improve performance of modern systems. Therefore understanding the contribution of out-of-order commit in the presence of prefetching is necessary. In gem5, we configure the L1-D cache with a stride prefetcher of degree 8. In Figure 10, we see the results of this study where hardware prefetchers and aggressive out-of-order commit perform similarly, with out-of-order commit providing slightly higher performance in all

Table 3: Out-of-Order Commit Costs for Unsafe Branches

Metric-PKI	SLM		NHM		HSW	
	Avg	Max	Avg	Max	Avg	Max
Branches	146.2	373.3	153.3	460.8	156.8	495.7
Mispred. br.	7.8	50.9	8.1	54.5	8.2	56.0
Num rollbacks	3.2	24.3	3.1	25.6	3.0	26.5
Instrs rollback	8.5	52.9	6.6	43.0	5.4	34.3

three architectures. When combining both out-of-order commit and prefetching in the SLM architecture, we see that it is almost additive, reaching to almost a 70% performance improvement (while prefetchers and out-of-order commit provide approximately a 40% speedup independently). In fact, combining aggressive out-of-order commit with prefetchers allows us to reach 84% of the ideal (sum of both) improvement for SLM, 77% for NHM, and 83% for HSW, demonstrating that these techniques work well together. This demonstrates that out-of-order commit can both work as an alternative to prefetching, as well as work together with it to provide a boost in single-core performance.

4.9 Rollback Costs for Unsafe Branches

While this work does not evaluate hardware costs, we are able to evaluate the number of rollbacks caused by committing past a mispredicted branch for each of the evaluated configurations. For this evaluation, we use a commit depth of 8, with aggressive out-of-order commit (AOOC) and Unsafe_BR (only the branch condition is not respected for out-of-order commit). In Table 3, we list the number of executed branches, mispredicted branches, the number of rollbacks caused by speculatively committing a mispredicted branch, and the number of instructions that were rolled back due to this rollback. While the number of executed branches increases with the aggressiveness of the core (because these cores can more aggressively speculate past branches), the average number of rollbacks (and instructions rolled back) decreases slightly with the aggressiveness of the core. These aggressive cores resolve speculative state quickly, reducing the number of instructions that need to be committed out-of-order. This results in a similar number of rollbacks per thousand architecturally committed instructions, around 3, for all configurations.

5 RELATED WORK

The goal of this work is to provide a detailed understanding into the potential for performance benefits across different out-of-order commit conditions and levels of aggressiveness. While this work aims to describe the maximum potential benefit for each individual condition, there have been many previous works that describe hardware solutions for early release of hardware structures and out-of-order commit strategies. Below, we provide an overview of these works, and how they fit into the categories described in this work.

5.1 Speculative Release of Hardware Structures

A number of implementations require register and processor state checkpointing support to speculatively retire or release hardware structures [1, 6, 7, 11, 14]. Processor state checkpointing, especially with the advent of very large SIMD registers such as AVX-512 registers which now support up to 32 registers with up to 512 bits per register, can require a significant amount of state to be saved when speculation is aggressive.




Figure 11: Maximum potential performance improvement is provided by Unsafe.OOC in which all conditions are unsafe. By preserving all conditions, maximum performance is reduced. This figure shows the negative effect of preserving (or making safe) a single condition compared to the maximum potential performance improvement. Safe.XX is equivalent to disabling all out-of-order commit conditions (all unsafe, highest performance) where XX indicates the only safe and preserved condition.

5.2 Non-speculative Structure Release

Non-speculative early release of hardware structures before commit requires knowledge that no older instruction (that has come earlier in the instruction stream) can cause the program to abort, raise an exception, or require exposure of the architected state at that time. Non-speculative solutions have the potential to be the most energy efficient, a necessity in an era of the end of Dennard Scaling for power-limited platforms. A range of solutions, from hardware-only to software-assisted solutions are described below.

Compiler support. Two previous works [1, 8] allow for early commit or reclaim of resources based on compiler knowledge, but require software recompilation.

No checkpoints. A number of solutions do not use checkpoints [3, 8, 13] and therefore do not require speculation (they respect the all commit conditions) or used software help to extend knowledge to the hardware.

Selective early commit. Early commit of loads [9, 10] allows for loads that have not yet received data from the memory hierarchy to become part of the committed state of a processor. This allows the processor to continue to process instructions past normally blocked structures, improving performance for memory-bound workloads. To accomplish this, the authors [10] decouple page faults that can occur from the fetching of data. While recompilation is not strictly required for this technique, the authors evaluate their work using a compilation strategy to expose loads early.

6 CONCLUSION

To obtain higher performance, extending the reach of the processor core has been a primary focus of much of microarchitecture research. One promising direction is the use of out-of-order commit, which releases precious processor resources early to allow the processor to extend its reach past typical hardware limits. In this work, we present a limit study for out-of-order commit through the introduction of reluctant and aggressive out-of-order commit modes. We show how smaller processors, even with a limited commit scan depth, can benefit from out-of-order commit strategies, but that larger, aggressive cores require deeper commit scan depths

to achieve improved performance. In addition, we provide a detailed breakdown of the contributions for each out-of-order commit condition for the SPEC CPU2006 benchmark suite. Our results show a very high potential for performance improvement, above 2.25x for some benchmarks, and believe that out-of-order commit strategies can play an important role for future energy-efficient and high-performance processor designs.

REFERENCES

- [1] F. Afram, H. Zeng, and K. Ghose. 2013. A Group-commit Mechanism for ROB-based Processors Implementing the x86 ISA. In *HPCA*. 47–58.
- [2] R. Badalane. 2015. DRAM’s Surprising Role in the Cost of Data Centers. <http://www.datacenterknowledge.com/archives/2015/11/12/dont-forget-memory-drams-surprising-role-high-cost-data-centers/>. (Nov. 2015).
- [3] G. B. Bell and M. H. Lipasti. 2004. Deconstructing Commit. In *ISPASS*. 68–77.
- [4] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill, and D. A. Wood. 2011. The gem5 Simulator. *SIGARCH Comput. Archit. News* 39, 2 (Aug. 2011), 1–7.
- [5] Intel Corporation. 2016. Intel® 64 and IA-32 Architectures Optimization Reference Manual. <http://www.intel.com/content/www/us/en/architecture-and-technology/64-ia-32-architectures-optimization-manual.html>. (June 2016).
- [6] A. Cristal, D. Ortega, J. Llosa, and M. Valero. 2004. Out-of-Order Commit Processors. In *HPCA*. 48–59.
- [7] A. Cristal, O. J. Santana, M. Valero, and J. F. Martínez. 2004. Toward Kilo-instruction Processors. *ACM Trans. Archit. Code Optim.* 1, 4 (Dec. 2004), 389–417.
- [8] N. Duong and A. V. Veidenbaum. 2013. Compiler-assisted, Selective Out-of-order Commit. *IEEE Computer Architecture Letters* 12, 1 (Jan. 2013), 21–24.
- [9] L. Gwennap. 1994. Digital Leads the Pack with 21164. In *Microprocessor Report*, 8(12). 249–260.
- [10] T. Ham, J. L. Aragón, and M. Martonosi. 2015. DeSC: Decoupled Supply-compute Communication Management for Heterogeneous Architectures. In *MICRO*. 191–203.
- [11] A. Hilton and A. Roth. 2010. BOLT: Energy-efficient Out-of-order Latency-tolerant Execution. In *HPCA*. 1–12.
- [12] D. Lee, Y. Kim, V. Seshadri, J. Liu, L. Subramanian, and O. Mutlu. 2013. Tiered-latency DRAM: A Low Latency and Low Cost DRAM Architecture. In *HPCA*. 615–626.
- [13] S. Marti, J. Borras, P. Rodriguez, R. Tena, and J. Marin. 2009. A Complexity-effective Out-of-order Retirement Microarchitecture. *IEEE Trans. Comput.* 58, 12 (2009), 1626–1639.
- [14] J. F. Martinez, J. Renau, M. C. Huang, and M. Prvulovic. 2002. Cherry: Checkpointed Early Resource Recycling in Out-of-order Microprocessors. In *MICRO*. 3–14.
- [15] J. E. Smith and A. R. Pleszkun. 1988. Implementing Precise Interrupts in Pipelined Processors. *IEEE Trans. Comput.* 37, 5 (May 1988), 562–573.
- [16] G. S. Sohi and S. Vajapeyam. 1987. Instruction Issue Logic for High-performance, Interruptable Pipelined Processors. In *ISCA*. 27–34.