

Improving the Performance and Power Efficiency of Shared Helpers in CMPs

Anahita Shayesteh
Comp Science Dept, UCLA
anahita@cs.ucla.edu

Glenn Reinman
Comp Science Dept, UCLA
reinman@cs.ucla.edu

Norm Jouppi
HP Labs, Palo Alto
norm.jouppi@hp.com

Tim Sherwood
Comp Science Dept, UCSB
sherwood@cs.ucsb.edu

Suleyman Sair
Electrical Eng Dept, NCSU
ssair@ncsu.edu

ABSTRACT

Technology scaling trends have forced designers to consider alternatives to deeply pipelining aggressive cores with large amounts of performance accelerating hardware. One alternative is a small, simple core that can be augmented with latency tolerant helpers. As the demands placed on the processor core varies between applications, and even between phases of an application, the benefit seen from any set of helpers will vary tremendously. If there is a single core, these auxiliary structures can be turned on and off dynamically to tune the energy/performance of the machine to the needs of the running application.

As more of the processor is broken down into helpers, and additional cores are added to a single chip that can potentially share helpers, the decisions that are made about these structures become increasingly important. In this paper we describe the need for methods that effectively manage these helpers. Our counter-based approach can dynamically turn off three helpers on average while staying within 2% of the performance when running with all helpers. In a multicore environment, our intelligent and flexible sharing of helper provides an average 24% speedup compared to static sharing in conjoined cores. Furthermore we show a benefit from constructively sharing helpers among multiple cores running the same application.

Categories and Subject Descriptors:

C.1 [Processor Architectures]: Multiple Data Stream Architectures

C.3 [Special-Purpose and Application-based Systems]: Microprocessor/microcomputer applications

General Terms: Design, Performance.

Keywords: CMP, factored core, helper engine, phase, helper

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CASES'06, October 23–25, 2006, Seoul, Korea.

Copyright 2006 ACM 1-59593-543-6/06/0010 ...\$5.00.

configuration, sharing policy, constructive sharing, flexible sharing.

1. INTRODUCTION AND MOTIVATION

While the unceasing march of Moore's law has given computer architects a continually increasing number of transistors to design with, emerging technology trends including poor wire latency scaling, increased power density, and reduced transistor reliability threaten to limit the usefulness of those designs. Chip Multi-processing based designs [5] address some of these scaling constraints by statically partitioning processor resources on the die into different cores that can exploit available parallelism in a more power-efficient manner.

Researchers have also proposed decoupling the large memory-laden structures that are crucial to single thread performance from the critical processor loop and redesigning them as specialized *helpers* [6, 7, 13, 1]. The essence of these helpers is that they are latency tolerant. They can provide an added performance advantage without impacting core scalability.

In some ways, helpers are analogous to power grid architectures that supply major urban areas with electricity. While coal or nuclear power plants typically supply the majority of power needs, during periods of peak power demand, auxiliary power plants (often powered by natural gas) are used to meet temporary load requirements. In the context of microarchitecture, we have a simple processor pipeline that can be augmented by auxiliary helpers when the need for a particular speculative technique is high. In our analogy to the power grid, the presumed goal of utility companies is to avoid power outages that can adversely impact normal electrical consumption patterns. Similarly, as microarchitecture desingers, we would like to optimize performance by finding a way to determine when helpers are needed for our processor pipeline so that they can be efficiently managed. This is achieved by taking advantage of the dynamic reconfigurability or polymorphism of helpers and allowing a core to adapt to changing applications, workloads, or phases. This is particularly critical in embedded devices designed for general purpose use. Here, die area can be efficiently utilized by sharing helper engines among multiple cores [10].

In order to resolve these optimization problems, several questions must first be addressed - which resources should be

shared; how should they be allocated; how can we efficiently manage their power? To answer these questions at runtime, we need a set of shared helper management policies that can adaptively allocate resources in a way that takes into consideration the needs of each executing workload.

In this paper, we present novel rule-oriented shared-helper management policies that are empirical based. We show that through the application of metrics based on a few simple-to-gather run-time statistics, processor resources can be allocated to a set of running programs in a way that is nearly optimal.

We begin by examining a single-threaded core with a variety of helpers. Our metrics predict which resources will be most valuable to an executing program so that inefficient helpers may be put into a low power state without impacting overall performance. By using a statistic as simple as the number of cache hits, we can pick the best overall configuration in a single try. Furthermore, we show how these techniques can be extended to help guide the sharing of resources between multiple cores on a chip. We observe how different helpers have different sharing requirements and demands that result in that prior work may not be able to find the best sharing combination. We have implemented a variety of different helpers and extended our simulator to support multiple executing processes so that a detailed treatment of the subject can be presented.

This paper makes the following contributions:

- Analysis of program behavior and resource requirements in an architecture with many major components decoupled via helpers.
- Proposal of an intelligent mechanism to dynamically tune helpers to the specific needs of the application for optimal power/performance.
- Extension of this mechanism to intelligent sharing of resources among multiple cores. Prior work has only considered whether or not simple sharing is possible. We demonstrate how different types of helpers may be shared differently and how sharing must be *flexible* to attain maximal performance.
- Investigation of constructive sharing when the same application is allowed to share fetch state across different phases of execution and different input sets.

The rest of this paper is organized as follows. In Section 2, prior work on helpers and CMP sharing is discussed. Simulation methodology can be found in Section 3. Section 4 begins with a description of the different helpers that we implemented and an analysis of the varying requirements for different applications for different helpers. We explore a mechanism to enable or disable helpers, and further extend this to managing shared helpers in a multicore environment in Section 5. We conclude in Section 6.

2. RELATED WORK

In [17], Smith proposes a processor implementation that consists of several distributed functional units, each fairly simple and with a very high frequency clock. These units communicate via point-to-point interconnections that have short transmission delays. He then describes how surrounding this simple core pipeline with *helpers* that perform speculative tasks off the critical path results in enhanced overall

performance. Since the helpers are off the critical path, they can use slower transistors to reduce static power consumption. This is also the motivation behind our factored design, where the speculative structures are reduced to a bare minimum size to support nearby ILP but they are duplicated in larger sizes outside of the critical path for extracting distant ILP. On a follow-up paper [6], Kim and Smith discuss the microarchitecture and ISA that implements this distributed processing paradigm, which utilizes hierarchical register files and a global register file to hold global state. Wang et al. introduce virtual multithreading to run helper threads on processors that do not have built-in hardware support for multithreading [22]. In Section 4, we detail the prior work for each individual helper when discussing the implementation of that helper.

Sharing some processor resources among cores in a CMP setting was first proposed by Dolbeau and Seznec [4]. Kumar et al. [10] also examine this idea and present a more thorough evaluation of sharing. Dolbeau and Seznec [4] propose the CASH architecture as an intermediate design point between CMP and SMT architectures. A typical CASH architecture shares caches, branch predictors, and division functional units between dynamically-scheduled cores. Kumar et al. [10] also propose resource sharing between adjacent cores of a chip multiprocessor to reduce die area with minimal impact on performance. They investigate the possible sharing of floating-point units, crossbar ports, instruction caches, and data caches and provide detailed analysis of area savings that each kind of sharing entails. Both [4] and [10] examine a round-robin based access model where a resource is allocated to a particular core every cycle. Kumar et al. also investigate a more sophisticated scheme for caches. After suffering a cache miss, a core relinquishes the control of the cache to the other core until the miss is serviced.

Our work differs from these techniques in two dimensions: 1) In addition to reactive (or demand-based) resource sharing such as when sharing caches, we also consider sharing always active resources such as prefetchers and value predictors that run ahead of the execution stream and, 2) We investigate an intelligent approach to assigning resources to cores based on utilization.

3. METHODOLOGY

The simulator used in this study was derived from the SimpleScalar/Alpha 3.0 tool set [3], a suite of functional and timing simulation tools for the Alpha AXP ISA. The timing simulator executes only user-level instructions. Simulation is execution-driven, including execution down any speculative path until the detection of a fault, TLB miss, or branch misprediction. Latency values for the caches and register files were obtained using CACTI [16] for a *70nm* process technology at 4 GHz.

We used the SPEC2000 benchmark set for our experiments. Although the results are gathered for all the benchmarks, we only show results for a randomly selected subset of 12 programs in the suite to conserve space in this paper. Details for all benchmarks will be available as a technical report. The programs were compiled on a DEC Alpha AXP-21164 processor using the DEC C and C++ compilers under OSF/1 V4.0 operating system with full compiler optimization (-O4 -ifo). We picked the 4 most dominant phases as determined by the hardware phase detection technique described in [15] and simulated these phases as representative

	Core	Helpers
Instruction Window and Physical RF	100 entry RF0	128 or 512 entry ROB 128 or 512 entry RF1
BBTB	256-entry 4-way set associative	1024-entry 4-way set associative
L1 Data Cache	8KB 4-way set associative, dual port with a 32 byte block size, 2 cycle latency	64KB 4-way set associative, single port with a 32 byte block size, 5 cycle latency
L1 Instruction Cache	4KB 4-way set associative, single port with a 32 byte block size, 2 cycle latency	64KB 4-way set associative, single port with a 32 byte block size, 4 cycle latency
Value Predictor (2 predictions per cycle)	none	2K-entry stride, 8K-entry L2 markov 2 read and 2 write ports
Address Predictor (1 predictions per cycle)	none	2K-entry stride, 8K-entry markov 1 read and 2 write ports
Stream Buffer	none	32-entry FA buffer
Branch Misprediction		15 cycles
Core Width	8-way issue, 8-way decode, 8-way commit	
Memory and L2 Cache	152 cycle memory latency, 2MB, 8-way set associative unified (instruction and data) cache with a 64 byte block size and 20 cycles latency	
Functional Units	8 integer ALUs, 2 integer MULT/DIV, 2 FP ALU, 2 FP MULT/DIV, 2 load/store	

Table 1: Simulation parameters for a single core architecture. These parameters remain the same for the multicore case except for the L2 cache which becomes a 4MB, 4-way set-associative, 4-bank cache shared among all cores.

samples of the program. On average, they accounted for approximately 70% of the execution time of each benchmark.

Table 1 presents the simulation parameters for the architecture we explore in this paper. We include an 8K entry gshare branch predictor in our model. In addition to modeling all of the structures and latencies in the architecture, we have extended SimpleScalar to include a cycle accurate, execution driven model of chip multiprocessing (CMP) [5]. All the parameters used in our multicore experiments are the same as in Table 1 for each core, except that we increase the size of the L2 cache to an 4MB, 4-way set-associative cache shared among all cores.

Per-thread performance metrics are measured for execution up to a maximum per-thread instruction count. All completed threads continue execution past this point while other threads execute. This prevents freeing of resources when certain threads complete earlier than others.

We make use of *weighted speedup* [18] as a performance measure. This metric ensures that high IPC threads are not favored and that we are measuring real increases in the rate of progress of all applications in the mix. Weighted speedup equalizes the contribution of each thread to the sum of total work completed in the interval by dividing the IPC of that job in the mix by the IPC of a single threaded run.

4. TUNING A SINGLE CORE

As our goal is to develop techniques to make good choices between a variety of different run-time processor configurations, we need to begin with a processor model that has been highly decoupled. We use the decoupled architecture proposed by Kursun et al. in [12]. Figure 1 illustrates the architecture we explore in this study, with many previously proposed helpers decoupled from the core pipeline. In this section we limit ourselves to single core designs, and then in Section 5 we show how to extend the ideas developed here to apply to cases where multiple cores may compete for resources.

4.1 Helpers

Because there is no existing simulation infrastructure with built-in decoupled helpers, we modified our CMP-extended version of SimpleScalar [3] to include a variety of proposed designs. The helpers we include, when combined together,

decouple most of the major pieces of the modern processor. This section includes a description of our target processor core and the helpers that assist it.

We consider a small L0 cache in our core, as in [7]. Our L1 data cache helper extends the cache hierarchy, providing larger capacity than the L0 at a faster latency than our L2. While data caches are perhaps the most well studied part of the processor, we will later show the benefits of controlling the data cache configuration in conjunction with other helpers.

Similar to the data cache, we make use of a smaller L0 instruction cache and an L1 instruction cache helper. To compensate for the smaller cache size, we use out-of-order instruction fetch as described in [20]. In this scheme, a placeholder is used in the instruction fetch queue (IFQ) to maintain program order - and the execution core stalls if the next entry to be consumed from the IFQ is still in flight. We model the complexity this brings to the IFQ by implementing the equivalent of an MSHR [8] for the instruction cache.

We model a stream buffer architecture guided by a stride-filtered markov predictor as proposed in [14]. The stream buffers of the prefetch helper are *only* accessed on L0 data cache misses. We allow a single prediction and a single prefetch per cycle, guided by the address predictor trained on the L0 miss stream.

Value prediction is one approach to break true data dependencies and create more instruction level parallelism in an application. We use a hybrid value predictor [21] to predict load instructions only. This structure can be accessed early in the pipeline as we only need the PC of the instruction to make the prediction. Our value predictor helper is limited to two predictions per cycle. We make use of an extra bit associated with each instruction in the instruction cache to dynamically mark instructions for value prediction. In this study, we only mark instructions that are loads.

Our architecture makes use of a basic block target buffer (BBTB) [23], a branch address predictor that predicts an entire basic block each cycle. The PC at the head of the basic block serves as an index to the predictor, which returns a target address, a fallthrough address, a branch type, and two per-branch prediction counters: one to make per-branch direction predictions, and one to arbitrate between

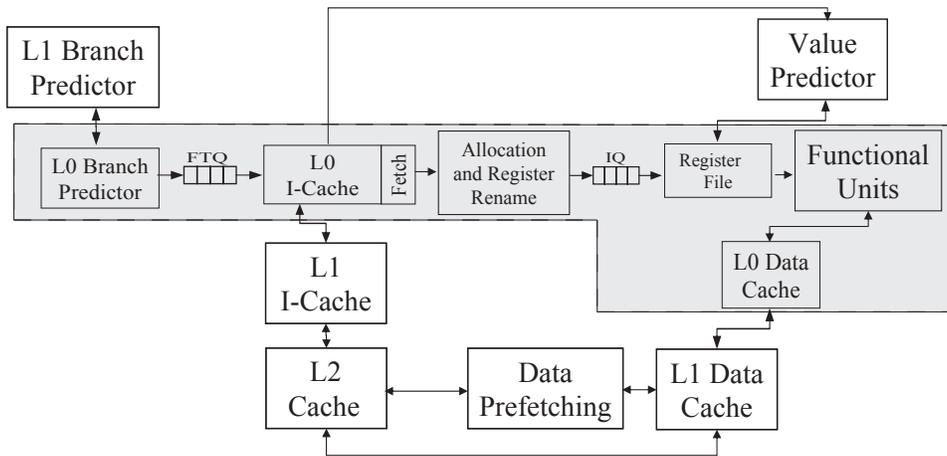


Figure 1: The single core architecture explored in this study.

the per-branch prediction and the global branch direction prediction.

Our architecture makes use of a small first level BBTB and a second level BBTB helper, similar to [13]. Similarly we decouple branch prediction from the instruction cache using a fetch target queue (FTQ) [13]. On a first level BBTB miss, the BBTB helper is probed and fetch stalls until a response is received from the helper. If the helper also misses, we guess a fixed fetch block size and continue fetching until a misprediction is detected.

4.2 Helper Utilization

While there are many circuit level advantages to decoupling large structures from the processor core, it also makes it very easy to *tune the processor to the specific needs of an application*. Each of these helpers has a well defined interface, and the processor can tolerate a wide range of access latencies to any one of these helpers (further explored in Section 5). Thus, adding supplemental control to each of these devices that will allow them to maintain separate power states does not introduce unreasonable cost or performance degradation.

As might be expected, for any given program some helpers will be more helpful than others. If the program is spending all of its time doing data accesses, it is more likely to get benefit from the data cache and prefetcher than the instruction cache. In a naive design, one might leave all of the helpers in an “on” state at all times. Clearly, this will be wasteful if a program gets no benefit from a subset of the helpers. In fact, we have found that up to half of the helpers can be turned off at any given point in time, with almost no performance impact (2%). The problem is knowing *which* helpers to turn off.

Note that there is an overhead associated with power gating, and the coarse-grain organization of these auxiliary structures into helpers provides a useful abstraction for power gating these structures. Moreover, phase-based optimization helps to hide the latency associated with power gating, since the granularity of our helper allocations is much larger than the latency of power gating.

To further illustrate this, we plotted the set of critical helpers that are needed to maintain top performance for each program in Figure 2. Figure 2 shows the minimal set

of helpers needed to achieve maximal performance for the top four phases of 11 of the programs we have examined (shown in the rows). The columns of the table represent the helpers, (d=data cache, p=prefetching, b=branch prediction, i=instruction cache, and v=value prediction). An X in a given square indicates for that phase, the corresponding helper should be turned “on”. Helpers without X’s can be disabled without affecting performance. We refer to setting the helpers “on” or “off” as the configuration of the helpers in one of two power states.

All of the configurations in Figure 2 perform within 5% of the configuration shown to have the absolute highest performance (typically the configuration with all helpers active). The last two columns present the speed up of this configuration relative to best and worst configurations (all helpers on, and all helpers off). These configurations were found strictly by a brute force search of the design space, simulating every possible configuration and taking the configuration with the least number of helpers “on” that was still within 5% of the case with all helpers turned on.

Obviously, trying each of the 2^n possible configurations is not a viable design choice for a runtime system, but there are several important points that can be drawn from this graph. It is clear from Figure 2 that there is no *one* good configuration that fits all applications. For many programs different configurations are even needed for different phases. This means a new, intelligent, and adaptive management scheme is going to be needed.

4.3 Helper Configuration

Helper configuration could be tackled in a variety of different ways. While static approaches could use profile or compiler guided heuristics to find one good configuration for the application as a whole, this may prove ineffective due to the time varying behavior of applications. A more effective static approach may be to inject special reconfiguration instructions that help tune the processor to specific phases. While these techniques are possible, in this paper we choose to focus on hardware based dynamic techniques as they require no a priori knowledge of the program and operate in a completely on-line manner.

To guide our decision about which configuration to choose, we need some information on how a program interacts with

the processor. This information must be simple and cheap to obtain, and should be highly indicative of the benefits the program is reaping from access to each helper. In order to gather this information, we use simple performance counters that track the “help” that each helper provides. For example, performance counters can track the number of hits in cache helpers, successful predictions by predictor helpers, and so forth. Specifically, our helpers are modified to track the following events:

Data/Instruction Cache: the number of times a cache line that missed in the L0 data/instruction cache hits in the data/instruction cache helper.

BBTB: the number of times a PC that missed in the L0 BBTB hits in the BBTB helper.

Data Prefetcher: the total number of hits in the stream buffers and in the L2 cache that were brought in by the prefetcher. Each line in the L2 cache is augmented with a bit to indicate whether it was brought in by a demand miss or a prefetch. On the first use of a line marked as a prefetch, the bit is flipped.

Value Prediction: the number of instructions issued using a predicted value.

These counter values are compared to pre-determined threshold values to decide whether the helper engine should be turned on or off. We performed a sensitivity analysis to various threshold values for the SPEC 2000 benchmark suite with reference inputs. We show results for a conservative set of thresholds to curtail performance degradation. More aggressive thresholds can be used when power reduction is the primary objective. Furthermore, dynamic threshold values can be calculated by adjusting the thresholds to maintain certain “on” and “off” state performance counter standard deviation values. However, this is beyond the scope of this paper, and we leave this for future work.

We make use of phase-based memoization [15] to track the helper configuration per application phase. A small hardware structure tracks a bit vector per phase for all the helpers in the architecture. If the bit at a particular location is set, the helper represented by that particular location should be on. Otherwise the helper can be turned off. The first time a phase is seen, we turn all helpers on and track the performance using the above counters. Each counter is compared against the threshold for keeping the helper on, and the bit vector for that particular phase is updated in our phase-based memoization table. The helpers can then be guided by a simple last phase predictor. The sampling period need not be as long as a phase, and can be limited to an interval of one million cycles.

We also memoize the observed IPC for each phase during the sampling period, and if the IPC resulting from a particular configuration is not within some threshold IPC seen during sampling (5% for this study), we clear the bit vector for that phase and force another sampling of the above counters. This helps recover from phase mispredictions or events that can impact performance (such as power throttling).

4.4 Performance of Counter-Guided Configuration

Figure 3 compares the performance of the configuration from Figure 2 (light grey) to our counter-guided configuration. Performance here is normalized to an architecture

running with all helpers on. On average, we are able to come within 1.5% of that performance with an average of only 2.6 helpers turned on. The configuration from Figure 2 comes within 2% of the performance of all helpers on, only using 2.2 helpers on average. In some cases, we use one more helper than needed due to our choice of conservative thresholds. More often than not this is the data cache helper, which can provide load hits that are not performance critical. One approach to fine tuning this further would be to incorporate the notion of load criticality [19] or to try and correlate this with the number of L2 misses – applications with a large number of L2 misses may not see benefit from an increase in L1 hits if the L2 misses that dominate the critical path of the application are not reduced.

These results demonstrate the ability of performance counters to fine tune helper engine utilization. This can allow an architecture to reduce power wasted in helpers that do not provide useful work or as we will see in the next section, to coordinate sharing among cores using a common pool of helpers.

5. MANAGEMENT ACROSS CORES

As demonstrated in Section 4, resource demand varies significantly across different applications and even across different phases of the same application. In the case of a single core, this fact can be exploited to reduce power by finding a configuration that still provides good performance but with a bare minimum of helpers left in a high power state. Managing helpers when multiple cores are involved presents a tougher challenge. While the easiest approach to supporting multiple cores on a chip would be to give each one its own set of helpers, previous work has shown that this instills unnecessary area complexity without a significant performance benefit [4, 10]. The alternative that we also explore in this paper is the use of a common pool of helpers, shared among all the cores. Sharing has a number of benefits such as reducing the area spent to implement redundant functionality and the potential of optimization through dynamic resource allocation. In this section, we present techniques to effectively manage helpers in a multicore environment.

5.1 Design Decisions

There are a large number of design decisions to be made in examining helpers in a multicore setting: the number of cores, the number of each type of helper, the topology of the interconnect between cores and helpers, the physical layout of cores and helpers, and the application mix to execute on the cores. This is an enormous design space, and it is simply not manageable to try all possible combinations. To get a set of experiments which is tractable, we limit our search in this paper by considering results for three possible machine organizations: two cores sharing a single helper of each type, four cores sharing a single helper of each type and four cores sharing two helpers of each type. Additionally, we consider applications that are not cooperative, but our work could certainly be applied to cooperative multithreading. We assume that all cores share a common second level cache, and that any core may connect to any helper.

5.1.1 Helper Latency Tolerance

One question that immediately comes to mind when we propose that any core may connect to any helper is that it is going to take a good deal longer to communicate with a

		d	p	b	l	v	S U/all	S U/none
applu	1		X				-4%	9%
	2						-2%	0%
	3		X				-4%	20%
	4		X		X		0%	19%
apsi	1			X	X		-3%	40%
	2						0%	2%
	3						-4%	0%
	4						-4%	0%
art	1		X			X	0%	16%
	2		X				-4%	8%
	3					X	-4%	7%
	4		X				-4%	8%
bzip2	1		X			X	0%	23%
	2					X	-2%	8%
	3		X				0%	23%
	4		X			X	-3%	10%
crafty	1	X		X	X		-4%	174%
	2	X		X	X		-3%	138%
	3	X		X	X		-2%	192%
	4	X		X	X		-4%	156%
eon	1	X		X	X		-2%	125%
	2	X		X	X		-1%	110%
	3	X		X	X		-2%	107%
	4	X		X	X		-1%	105%

		d	p	b	l	v	S U/all	S U/none
equake	1		X				-1%	79%
	2			X		X	-1%	122%
	3					X	-3%	3%
	4			X		X	-2%	112%
galgel	1						-4%	0%
	2						-1%	0%
	3						-1%	0%
	4						-3%	0%
gap	1			X	X	X	-3%	174%
	2			X	X		-4%	40%
	3			X	X		-4%	36%
	4		X			X	0%	85%
mcf	1		X			X	-1%	985%
	2					X	-2%	15%
	3					X	-3%	407%
	4					X	-2%	10%
mesa	1			X	X	X	-1%	109%
	2			X	X	X	-1%	94%
	3			X	X	X	-1%	101%
	4			X	X	X	-1%	102%
parser	1	X		X		X	-2%	29%
	2	X		X		X	-1%	24%
	3	X		X		X	-4%	17%
	4	X		X		X	-1%	18%

Figure 2: Optimal helper configurations for top executed phases of benchmarks

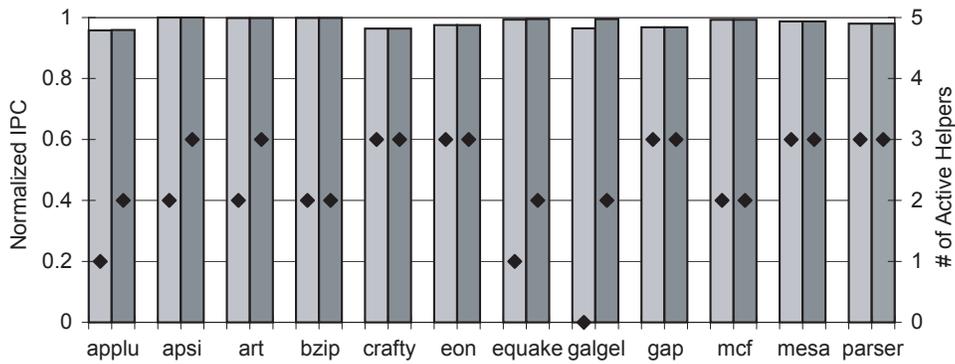


Figure 3: Comparison of the configuration from Figure 2 (light grey) and our counter guided configuration (dark grey).

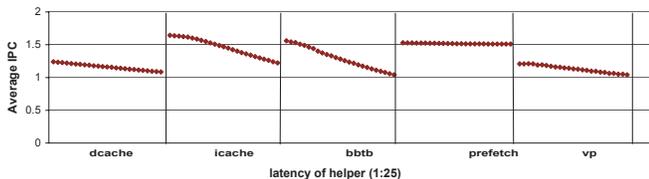


Figure 4: Impact of Latency on Our Helper Engines

helper on the other side of the chip than with one in close physical proximity to the core. While this is true when partitioning arbitrary processor resources, helpers have an inherent advantage due to their higher latency tolerance.

To demonstrate that helpers are naturally latency tolerant, we present Figure 4 which plots the performance impact of latency on our various helpers by varying the access latency from 1-25 cycles. For these results, we consider a single core with private helpers, and average the IPC observed over benchmarks that use the helper from table 2. As seen from the figure, for most helpers, there is little impact on performance of these helpers from smaller latencies – the most is seen by the branch target buffer (BBTB) helper, which suffers an IPC degradation around 1% for each additional cycle of latency. Prefetching sees the least impact, less than a 0.01% drop in IPC for each additional cycle of latency. The prefetcher hides the latency of memory, and even 25 cycles is tolerable when compared with this latency. The remaining helpers see less than 0.5% degradation per cycle for each additional cycle. As the latency increases above 10 cycles, its impact on the performance of the helper increases non-linearly.

Our architecture is also not impacted by nonuniform access latency from different cores to a common helper. The arbiter that selects what requests from a core should be serviced by a helper would be located close to the helper itself. Therefore, if core A sees a two cycle latency to a helper and core B sees a single cycle latency to a helper, and if core A pipelines its requests over two cycles, then the helper will simply see a stream of requests from A and B without any notion of heterogeneous latency. A will see its predictions a cycle later than B will see its predictions, but as we have demonstrated, this has a negligible impact on performance.

Our multicore architecture and its floorplan are shown in Figure 5. We use CACTI to estimate the size and dimensions of the L2 cache and all of our helpers. The area of our core was calculated using the area of EV6 and EV5 scaled to a 70nm feature size, similar to [9].

Our flexible sharing requires a link from each core to each helper. This would double the number of interconnections compared to a conjoined architecture where each helper is statically shared between two cores. We used a similar method of crossbar area estimation as [10]. For our choice of helpers and their respective bandwidth requirements, the crossbar area occupies 10% of the total area of the processor, which is 5% more than conjoined cores.

The other hidden cost in sharing helpers is the potential increase in requests for each resource, making helper bandwidth a serious concern. Ideally, each core would have its own dedicated port to each helper, but the cost would be prohibitive. Instead, a helper can make use of port arbitration to satisfy multiple core requests. One possibility is al-

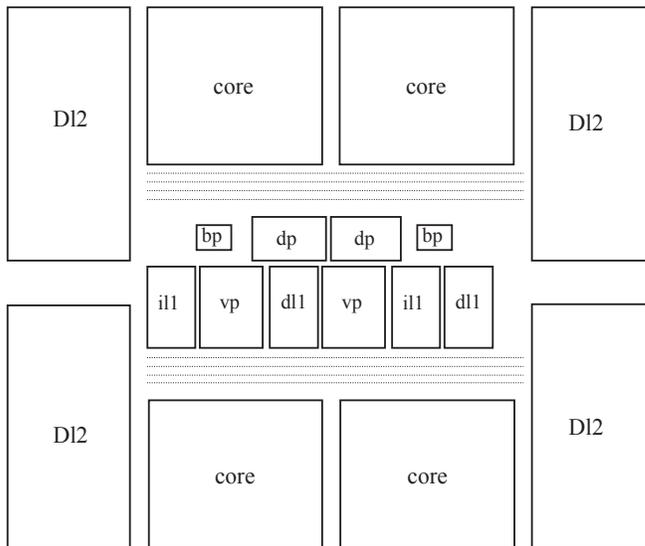


Figure 5: Our CMP floorplan, with cores and L2 cache banks distributed around 2 sets of helpers.

lowing cores to take turns accessing a helper. Another would involve more sophisticated control hardware that would arbitrate among several requests, much like what is done with a unified second level cache.

In an architecture where cores share a common pool of helpers, the helpers can either be exclusively assigned to a core (i.e. partition the helpers), shared among several cores (i.e. sharing common helpers), or some combination of both (i.e. some of the cores sharing a common helper). Partitioning is useful when the bandwidth or internal storage demands placed on a helper by a single core would preclude benign sharing with other cores. In that case, the most favorable option would be dedicating a helper to a single core. Sharing can be useful when individual cores do not consume all of the bandwidth or storage space of a given helper.

5.1.2 Always On vs. On Demand Sharing

Different helpers exhibit different tolerances to sharing. At a high level, helpers can be divided into those that are accessed on-demand and those that are always active. On-demand helpers include those that are hierarchical extensions of core structures, like the instruction cache, data cache, and branch predictor. These helpers are only accessed when their corresponding core structure misses. Locality in the corresponding core structures filters the majority of the requests obviating frequent accesses to these helpers. This means that bandwidth to the helper may more easily be shared among multiple cores. However, the amount of state contained in the helper may still be insufficient to allow effective sharing between cores – but bandwidth is not usually a limiting factor.

Other helpers do not have corresponding core structures, and are therefore not on-demand. The value prediction and prefetching helpers are examples of this class of helpers. Any load instruction can be value predicted and any cache miss can initiate a prefetch stream, but a helper may not have enough bandwidth to handle competing requests from multiple cores if there are lots of loads or cache misses. Helper

state is also a problem here, as sharing cores would need to contend with one another for value and address predictor space.

5.1.3 Simple vs. Counter Guided Sharing

One simple approach to sharing is to have conjoined cores take turns accessing a common set of helpers. Taking turns is a viable option for on-demand helpers where contention for helper bandwidth is less common. But helpers that are not on-demand can suffer from a turn-based approach. Consider value prediction in a two core setting, where each cycle a new set of load PCs are fetched and could potentially be value predicted. Taking turns would mean that value prediction would only occur every other cycle, potentially halving the number of value predicted instructions. Note that the issue here is not access latency (because helpers are latency tolerant), but the fact that we lose the opportunity to predict load instructions. This problem only becomes worse when trying to share among more than two cores. We further optimize the turn-based strategy by allowing other cores to access a helper on a given core's turn if that core does not require the helper's bandwidth. This can happen when a core has suffered a pipeline flush or in the case of on-demand helpers that simply do not see any accesses to the helper.

Even though some cores may make use of a helper, the helper may not provide any benefit, even in the case of on-demand helpers. Consider a thread whose working set does not fit in the data cache helper – it will thrash in the data cache as it tries to contain its working set, but will still be plagued by misses that must be serviced by the L2 and will evict potentially useful entries from other threads. Similarly, a thread may not see any benefit from value prediction and may simply be stealing available bandwidth from a thread that does see benefit.

We make use of the utilization counters from Section 4 to guide helper sharing. These counters provide a filtering mechanism to avoid sharing a helper among cores that see no benefit from that helper. Cores can then request access from a global helper arbiter to the helpers from which they expect to see benefit.

While filtering useless sharing is vital, we also need to provide an intelligent approach to choose what cores will share a common helper. In a four core scenario where all cores (labeled A-D) want a BBTB helper, and there are only two helpers available, performance may substantially improve if A and B are allowed to share instead of A and C. Or, it may be best for A to have its own private helper and for B-D to share the remaining helper.

To arbitrate sharing among the filtered set of helpers, we make use of correlating counters that can guide sharing. A good example of this is the prefetching helper. If multiple cores are contending for a pool of prefetch helpers, threads that have a greater magnitude of prefetches should be given private access to a prefetcher if one is available. If there are not enough prefetchers to grant a private to threads with a large number of prefetches, threads with a comparable number of prefetches should be paired together. This prevents threads with a greater number of prefetches from starving other threads from getting access to stream buffers.

At each core, the correlating counters are tested against a number of thresholds to classify the demand for a particular helper into a utilization class (i.e. light, medium,

heavy). The utilization class is then memoized in the phase-detection hardware to track the expected utilization class of each helper at each phase.

The global helper arbiter is responsible for taking the requests from all cores for helpers (this has already been filtered at each core by the performance counters and phase-detection mechanisms) and the actual utilization class for each helper requested from each core. The global arbiter then tries to match requests of similar utilization classes together in the case of helpers that are not on-demand and tries to mix requests in the case of helpers that are on-demand. The global arbiter makes decisions about assignments of helpers to cores every 10 million cycles, the same granularity as our phase detection intervals. This granularity of decision making can easily absorb both the latency to communicate the counter values to the arbiter and the latency for the arbiter to make a decision based on these counters.

We will refer to the architecture that can allow helper sharing between any set of cores on a CMP our flexible sharing architecture.

5.2 Sharing Results

Utilization counters can potentially improve multi-threaded workload in two ways: filtering useless sharing and providing information to choose the best pairing of applications to share a helper. To better evaluate the contribution of each of these factors we study both 2-core and 4-core CMPs sharing helpers. The former case can only see benefit from filtering useless accesses as there is only one way to share helpers. The latter case can take advantage of flexible sharing of helpers as well as filtering useless sharing.

5.2.1 2-Core Results

Figure 6 compares the performance of our counter guided approach to a baseline that simply shares each helper among two cores. Utilization counters prevent applications from accessing the helper when they do not see benefit from that helper. Bars represent normalized IPC, and dots represent the reduction in the total number of accesses to the helpers. The application mixes are constructed to include threads with disjoint demand for helpers. For example, **apsi** and **crafty** heavily use instruction and branch prediction helpers (on-demand helpers) while **quake**, **gap**, and **mcf** see significant improvement using value prediction and prefetch helpers.

On-demand helpers do not see much improvement from counter guided approach because they already inherently filter requests as can be seen from small reduction in the number of accesses to the helper. However filtering useless accesses is essential to always-on helpers. **Quake** sees a 27% speedup just from eliminating useless accesses to the prefetch helper from **apsi**. The large reduction in accesses to the helpers for these mixes is another indication that the counter guided approach can improve performance and reduce power at the same time.

5.2.2 4-Core Results

Next, we examine the benefits of flexible sharing. First, to better evaluate the difference between the two types of helpers (section 5.1.2), we more closely examine the BBTB helper (an on-demand helper) and the prefetcher helper (not on-demand). For each, we examine situations where four

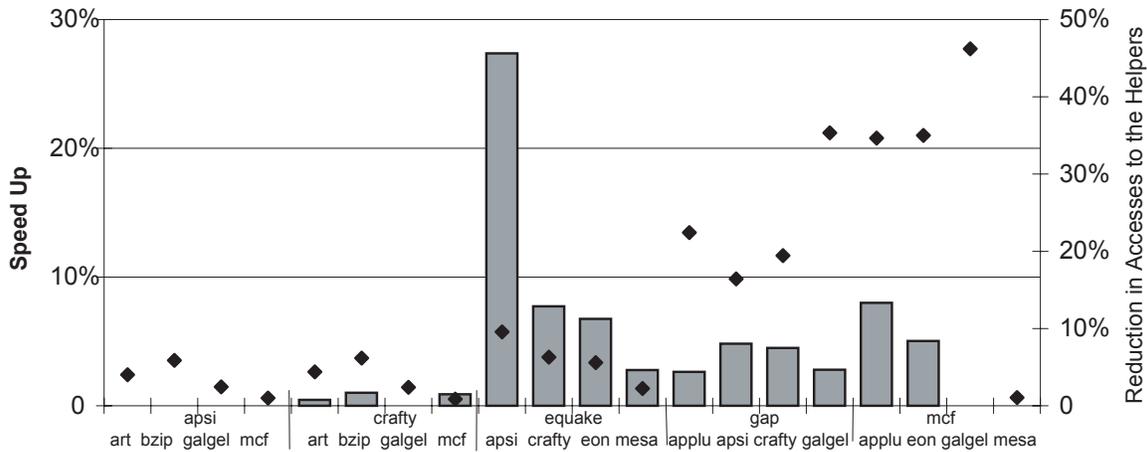


Figure 6: Impact of filtering useless accesses to a single set of helpers shared by two cores. Results shown are relative to conjoined sharing without filtering.

cores have private versions of all other helpers but share either the BBTB or the prefetcher. The four cores are then forced to share either one or two instances of either the BBTB or the prefetcher. In this section, with the exception of our baseline architecture, simple sharing *always* uses counter-guided filtering to reduce accesses to helpers. But cores may or may not be able to flexibly share helpers. In simple sharing schemes, cores may only share helpers with certain other cores.

We consider the following architectures:

baseline - Four cores share a single helper - this configuration does not use counter-guided filtering.

allprivate - All cores have their own helpers - no sharing is done.

1helper - Four cores share a single helper using counter-guided filtering.

average, worst, best - This configuration shares two sets of helpers among four cores. When using simple sharing, performance depends on what cores are actually sharing (i.e. conjoined). With four cores, there are three possible pairings. For simplicity, we show values for the best of the three configurations, the worst of the three, and the average over all three.

flexible - The flexible helper sharing policy from section 5.1.3. One core can share a prefetcher helper with another core, and a value predictor helper with a different core.

For each helper that we explore, we construct an application mix by selecting five benchmarks that benefited from the helper and five benchmarks that did not benefit from the helper. We then form five application mixes of four threads each that represent all possible combinations: all threads need the helper, three out of four threads need the helper, two out of four threads need the helper, only one thread needs the helper and no thread needs the helper - in that order. For cases where only one or no application demands a helper, there is obviously little impact from these sharing approaches - but these results are shown for completeness.

Our helper configurations are as presented in Section 3. The value predictor helper has only two ports, and can therefore only satisfy two requests for prediction per cycle. The prefetcher can only prefetch one cache line per cycle.

Figure 7 illustrates the performance of an on-demand helper: the BBTB helper. For the single helper runs, there is not much improvement from counter-guided sharing because on-demand helpers naturally will not be accessed if there is no benefit. However, when sharing two helpers among four cores, there is a destructive relationship between **eon** and **crafty** when sharing a common BBTB due to aliasing. Our counter-guided approach is able to determine the best combination for benign sharing - even when all four applications want to share the BBTB. Because **eon** and **crafty** see many more helper BBTB hits than other applications, they should not be combined together to avoid contention for space in the BBTB. The gap in performance between having a single BBTB helper and having two BBTB helpers is clearly greater when more applications demand the BBTB, demonstrating the impact that contention for space can have, even for on-demand helpers. The last two mixes of benchmarks, where only one or no application needs the BBTB, do not see any impact from sharing approaches.

Figure 8 illustrates the performance of a helper that is not on-demand: the prefetch helper. Counter-guided sharing is useful for the single helper run when there are two or three benchmarks competing for bandwidth that do not see any benefit from the helper.

The benchmark mix **mesa-parser-applu-quake** is such a case, where **mesa** and **parser** do not benefit from prefetching but **applu** and **quake** do. By filtering **mesa** and **parser** out via our correlation counters, we are able to use a single helper to outperform the worst and average cases of conjoined cores with two helpers.

With two helpers, there is more disparity between different conjoined core runs. **Applu** and **mcf** issue far more prefetches than **bzip2** and **quake**, and unlike the case of on-demand helpers, it is actually beneficial here to have **applu** and **mcf** share the same helper. If we combine either one of these applications with one that benefits from prefetching but does not have the same magnitude of prefetches (like **bzip2** or **quake**), the application with less prefetches

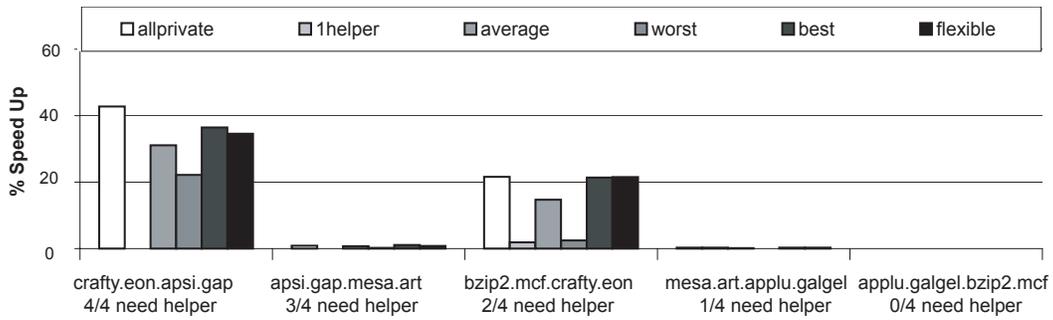


Figure 7: Simple vs Counter-Guided Sharing of the BBTB Helper Engine

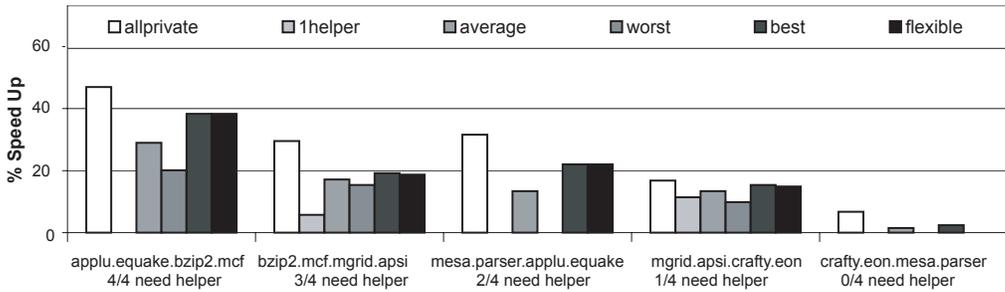


Figure 8: Simple vs Counter-Guided Sharing of the Prefetching Helper Engine

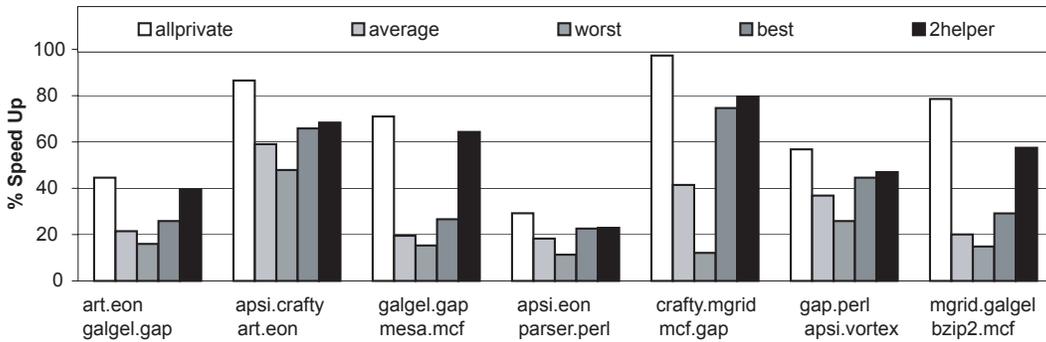


Figure 9: Simple vs Counter-Guided Sharing of all Helper Engines

will not get a fair share of stream buffer resources. In the case of `mesa-parser-applu-equake`, `parser` contends for prefetch bandwidth, despite not seeing a benefit from prefetching, and impedes the prefetching of `applu` or `equake`. The counter-guided approach is able to identify the inability of `parser` to effectively use prefetching, and prevents this degradation. For all runs, our counter-guided approach is able to perform as well or better than the best conjoined combination.

Figure 9 presents results for sharing all helpers at once. The first application mix on the figure enjoys a large improvement from our counter-guided approach for two main reasons. First, `gap` is able to get a private value predictor. Second, `art` and `eon` do not do well when conjoined with `galgel`. By giving flexibility to helpers in how they share, our approach is able to outperform any conjoined combination. This is also evident from the third benchmark mix, where value prediction contention hampers the performance of `mcf` and `gap`. On average our counter-guided approach for one set of helpers sees 13% improvement over a baseline naively sharing helpers among all cores. Sharing two set of helpers, our approach provides 54% improvement while conjoining cores can see benefit ranging from 20% to 40% depending on how cores are conjoined.

5.2.3 Constructive Sharing

As we demonstrated in the previous section, sharing one helper among four cores can degrade performance significantly when applications running on all cores need the helper. However, there are cases where common code or data (i.e. OLTP, parallel processing) may be executing in the CMP environment, and sharing can actually be constructive if threads are allowed to share state in a common helper. Our flexible helper management allows constructive sharing of a helper among such workloads. We consider the case where the same application is executing on all cores, but each application instance is executing different input sets or different phases of the same input set.

Figure 10 illustrates the benefit of constructive sharing for the BBTB and instruction cache helpers. For `crafty`, `eon` and `mesa` we simulated four different phases of each application running concurrently on four cores. For `vortex` we used four different inputs. All of these application use both BBTB and instruction cache helpers intensively. The first bar shows results when there are private helpers dedicated to each core, and the second bar shows results when only one helper is shared constructively among the four cores. The speedup presented is relative to one helper shared among four cores, but without any constructive sharing among threads. Our counter mechanism is still used in this case. Our results indicate that sharing the instruction cache between multiple cores incurs a similar miss rate as a dedicated cache per core with the same capacity, as reported in [11]. In `vortex`, a constructively shared instruction cache even outperforms the private cache performance by avoiding misses to cache blocks used by multiple cores. The contention when sharing the BBTB among applications that have high demand for this helper can be extremely severe when not shared constructively. `Crafty` and `vortex` see a 4X speedup when using dedicated BBTB helpers instead of sharing a single BBTB helper. This is due to the reduced accuracy of branch prediction when threads thrash for space in the shared predictor. There is a significant increase in the number of instructions

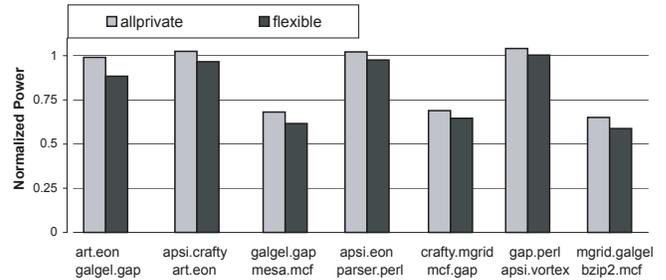


Figure 11: Power dissipated in flexible and allprivate configurations normalized to baseline configuration with shared helpers.

executed as a result of this mispeculation which can further pollute caches and waste energy. Sharing the BBTB constructively among cores eliminates this thrashing effect for all applications simulated except `crafty` which still sees some impact from this. This application has more complex branch behavior that can inhibit constructive sharing across different phases.

5.2.4 Power Savings from Flexible Sharing

To better understand the tradeoffs between sharing helpers and keeping private helpers for each core, we integrated Wattch [2] into our simulation infrastructure and obtained energy results for our flexible sharing policy and the configuration where every core has a private helper. We used process parameters for a 70nm process at 5.6GHz with 1V supply voltage. Our results are extracted with the most aggressive conditional clocking strategy, where dynamic power scales linearly with access to the ports.

Figure 11 illustrates that on average, the flexible sharing policy is able to reduce the energy dissipated, despite the fact that the performance of the flexible policy can be slightly worse than the all private configuration.

6. SUMMARY

In this paper we explore helper management policies for both single and multicore configurations. Cores with decoupled helpers provide an opportunity to dynamically tune processor resources on an application phase basis. With our counter-guided helper management, we can come within 2% of the best performing helper configuration (typically the one with all five helpers active) with an average of less than three helpers turned on. This counter-guided scheme can be applied to the multicore environment to more effectively share a pool of common helpers among a number of cores. Our approach to sharing is intelligent and flexible enough when used with four cores sharing two sets of helpers to see an average 54% weighted speedup over a baseline naively sharing only one set of helpers. Statically shared helpers can see benefit performance by an average of 20% to 40% depending on how cores are shared. Constructive sharing can provide even more benefit, effectively providing performance comparable to private helpers when running the same application on all cores, even for different inputs and phases.

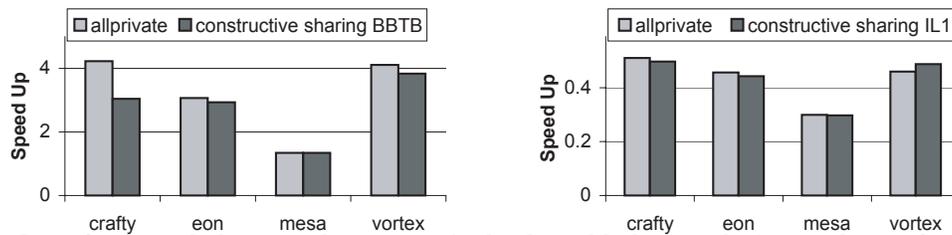


Figure 10: Speedup of private helpers and constructively shared helpers relative to the performance of sharing helpers destructively.

7. REFERENCES

- [1] R. Balasubramonian, S. Dwarkadas, and D. Albonesi. Reducing the complexity of the register file in dynamic superscalar processors. In *Proceedings of the 34th Annual International Symposium on Microarchitecture*, December 2001.
- [2] D. Brooks, V. Tiwari, and M. Martonosi. Wattch: A framework for architectural-level power analysis and optimizations. In *27th Annual International Symposium on Computer Architecture*, June 2000.
- [3] D. C. Burger and T. M. Austin. The simplescalar tool set, version 2.0. Technical Report CS-TR-97-1342, U. of Wisconsin, Madison, June 1997.
- [4] R. Dolbeau and A. Sez nec. Cash: Revisiting hardware sharing in single-chip parallel processor. Technical Report IRISA Report 1491, IRISA, November 2002.
- [5] L. Hammond, B. A. Nayfeh, and K. Olukotun. A single-chip multiprocessor. *IEEE Computer*, 30, 1997.
- [6] H-S. Kim and J. E. Smith. An instruction set and microarchitecture for instruction level distributed processing. In *Proceedings of the 29th annual international symposium on Computer architecture*, pages 71–81, June 2002.
- [7] J. Kin, M. Gupta, and W. Mangione-Smith. The filter cache: an energy efficient memory structure. In *IEEE International Symposium on Microarchitecture*, December 1997.
- [8] D. Kroft. Lockup-free instruction fetch/prefetch cache organization. In *8th Annual International Symposium of Computer Architecture*, pages 81–87, May 1981.
- [9] R. Kumar, N. Jouppi, P. Ranganathan, and D. Tullsen. Single-isa heterogeneous multi-core architectures: The potential for processor power reduction. In *36th International Symposium on Microarchitecture*, December 2003.
- [10] R. Kumar, N. Jouppi, and D. Tullsen. Conjoined-core chip multiprocessing. In *37th International Symposium on Microarchitecture*, December 2004.
- [11] P. Kundu, M. Annavaram, T. Diep, and J. Shen. A case for shared instruction cache on chip multiprocessors running oltp. *SIGARCH Comput. Archit. News*, 32(3):11–18, 2004.
- [12] E. Kursun, A. Shayesteh, S. Sair, T. Sherwood, and G. Reinman. An evaluation of deeply decoupled cores. In *Journal of Instruction Level Parallelism*, volume 8, 2006.
- [13] G. Reinman, T. Austin, and B. Calder. A scalable front-end architecture for fast instruction delivery. In *26th Annual International Symposium on Computer Architecture*, May 1999.
- [14] T. Sherwood, S. Sair, and B. Calder. Predictor-directed stream buffers. In *33rd International Symposium on Microarchitecture*, December 2000.
- [15] T. Sherwood, S. Sair, and B. Calder. Phase tracking and prediction. In *30th Annual International Symposium on Computer Architecture*, June 2003.
- [16] P. Shivakumar and Norman P. Jouppi. Cacti 3.0: An integrated cache timing, power, and area model. In *Technical Report*, 2001.
- [17] J. E. Smith. Instruction-level distributed processing. *IEEE Computer*, 34(4):59–65, April 2001.
- [18] A. Snaveley and D. M. Tullsen. Symbiotic jobscheduling for a simultaneous multithreading processor. In *In Ninth International Conference on Architectural Support for Programming Languages and Operating Systems*, November 2000.
- [19] S. Srinivasan, R. Ju, A. R. Lebeck, and C. Wilkerson. Locality vs. criticality. In *28th Annual International Symposium on Computer Architecture*, June 2001.
- [20] J. Stark, P. Racunas, and Y. N. Patt. Reducing the performance impact of instruction cache misses by writing instructions into the reservation stations out-of-order. In *30th International Symposium on Microarchitecture*, pages 34–43, December 1997.
- [21] K. Wang and M. Franklin. Highly accurate data value prediction using hybrid predictors. In *30th Annual International Symposium on Microarchitecture*, pages 281–290, December 1997.
- [22] P. Wang, J. Collins, H. Wang, D. Kim, B. Greene, K. Chan, A. Yunus, T. Sych, S. Moore, and J. Shen. Helper threads via virtual multithreading on an experimental itanium[®] 2 processor-based platform. In *ASPLOS-XI: Proceedings of the 11th international conference on Architectural support for programming languages and operating systems*, pages 144–155, New York, NY, USA, 2004. ACM Press.
- [23] T. Yeh and Y. Patt. A comprehensive instruction fetch mechanism for a processor supporting speculative execution. In *Proceedings of the 25th Annual International Symposium on Microarchitecture*, pages 129–139, December 1992.