

Improving Application Performance by Dynamically Balancing Speed and Complexity in a GALS Microprocessor*

Greg Semeraro¹, David Albonesi², Steven Dropsho³, Grigorios Magklis³, and Michael L. Scott³

¹ Dept. of Computer Engg.
Rochester Institute of
Technology
Rochester, NY 14623
gpseec@ce.rit.edu

² Dept. of Electrical and
Computer Engineering
University of Rochester
Rochester, NY 14627
albonesi@ece.rochester.edu

³ Dept. of Computer Science
University of Rochester
Rochester, NY 14627
{dropsho, maglis, scott}
@cs.rochester.edu

ABSTRACT

Microprocessors are traditionally designed to provide “best overall” performance across a wide range of applications and operating environments. Several groups have proposed hardware techniques that save energy by “downsizing” hardware resources that are underutilized by particular applications. We explore the converse: “upsizing” hardware resources in order to improve performance relative to an aggressively clocked baseline processor. Our proposal depends critically on the ability to change frequencies independently in separate domains of a globally asynchronous, locally synchronous (GALS) microprocessor.

We use a variant of our multiple clock domain (MCD) processor, with four independently clocked domains. Each domain is streamlined with modest hardware structures for very high clock frequency. Key structures can then be upsized on demand to exploit more distant parallelism, improve branch prediction, or increase cache capacity. Although doing so requires decreasing the associated domain frequency, other domain frequencies are unaffected. Measuring across a broad suite of application benchmarks, we find that configuring just once per application increases performance by an average of 17.6% *compared to the best fully synchronous design*. When adapting to application phases, performance improves by over 20%.

1. INTRODUCTION

*This work was supported in part by NSF grants CCR-9701915, CCR-9811929, CCR-9988361, EIA-0080124, and CCR-0204344; by DARPA/ITO under AFRL contract F29601-00-K-0182; by an IBM Faculty Partnership Award; and by equipment grants from IBM and Intel.

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.

Workshop on Application Specific Processors San Diego, CA USA

Microprocessor design traditionally embodies a tradeoff between processor frequency and the hardware complexity required to exploit instruction-level parallelism (ILP). Any particular design is typically a compromise, because the tradeoff is highly application dependent. Optimal pipeline depth, for example, tends to vary across classes of applications, *e.g.* SPEC versus transaction processing [13, 19, 28]. Studies have even demonstrated significant variability in hardware requirements among the different phases of a given application [8, 27, 30]. Any processor designed for “best overall” performance inevitably makes inefficient use of its hardware resources when running a varied workload. This observation suggests the possibility of dynamically optimizing the tradeoff between clock rate and IPC (instructions per cycle) [1].

The simplest approach is to “downsize” underutilized structures in order to save energy. Proposals to date include adaptive cache hierarchies [2, 8, 23], adaptive issue queues [5, 11], and combinations including caches, issue queues, register files, and the reorder buffer [9, 22].

A downsized structure often operates at a higher speed. For instance, Buyuktosunoglu [5] demonstrated a 70% reduction in issue queue access time when downsizing from 32 entries to eight. Unfortunately, critical paths elsewhere in the chip typically prevent this speedup from translating into an overall increase in clock speed.

If one aims to improve performance, an alternative strategy [1] is to *decrease* global frequency when one or more structures are *upsized*. This strategy preserves the correctness of critical paths, but succeeds only when the performance gained by upsizing some subset of the chip’s hardware resources (*e.g.*, the data cache hierarchy) exceeds the performance lost by slowing down everything else, and this is seldom the case.

Globally-Asynchronous, Locally-Synchronous (GALS) microprocessors [15, 20, 26] offer a new opportunity to dynamically trade speed for complexity within a *subset* of the processor. In our Multiple Clock Domain (MCD) GALS design [26], the chip is broken down into separate domains for front-end fetch / rename, integer execution, floating point execution, and load / store (L1 data cache and L2 cache). Clock frequency can be varied in each domain independent of the others; synchronization circuits are used on all cross-

domain paths. By varying voltage with frequency, we been able to save significant energy with only a modest loss in performance [17, 25, 26]. We use a similar architecture in the current paper in a quest to *improve* performance.

We begin with a baseline MCD processor in which each domain is optimized for high frequency, with relatively low hardware complexity. Dynamic upsizing of key hardware structures (with an accompanying decrease in domain frequency) then offers per-application tailoring of the speed / complexity tradeoff. In effect, we trade speed for complexity whenever an increase in potential ILP provides IPC benefits that override the frequency loss *in that one domain*. Our baseline processor has a higher branch mispredict penalty than found in a fully synchronous design, due to pipeline inefficiencies at lower frequencies. It also suffers a modest frequency penalty due to the ability to adapt. Despite these handicaps, we demonstrate overall performance improvements of 17.6% with respect to the *best fully synchronous design* using only profile-based whole-program adaptation of the instruction cache, branch predictor, integer issue queue, floating-point issue queue, and data / L2 caches. When we add an adaptation algorithm that automatically detects phases of program behavior and sizes structures appropriately for each phase, average performance improvement reaches 20.4%. These results reflect an exhaustive exploration of the space of synchronous and whole-program adaptive processor configurations, embodying approximately 300 CPU months of simulation time.

2. MCD MICROARCHITECTURE

Our MCD architecture has four independent clock domains, comprising the front end (L1 instruction cache, branch prediction, rename, reorder buffer and dispatch); integer processing core (issue queue, register file and execution units); floating-point processing core (issue queue, register file and execution units); and load / store unit (load / store queue, L1 data cache and unified L2 cache). The main memory can be thought of as a separate fifth domain, but is fixed at a base frequency and hence non-adaptive. Further details, including a description of the inter-domain synchronization circuitry, can be found in prior papers [17, 25, 26].

We investigated two models for dynamic voltage and frequency scaling: an *XScale* model and a *Transmeta* model, both of which are based on published information from their respective companies [7, 10]. For both of these models, we assume that the frequency change can be initiated immediately when transitioning to a lower frequency and voltage, while the desired voltage must be reached first before increasing frequency. For the Transmeta model, we assume a total of 32 separate voltage steps, at 28.6mV intervals, with a voltage adjustment time of 20 μ s per step. Frequency changes require the PLL to re-lock. Until it does, the domain remains idle. The PLL locking circuit is assumed to require a lock time that is normally distributed with a mean time of 15 μ s and a range of 10–20 μ s. With the tight interaction between the domains, the suspension of one domain quickly cascades to other domains. As one can imagine, this has a profound impact on overall performance. For these reasons the Transmeta model was not investigated further and is not included in any of the analysis that follows. For the XScale model, we assume that frequency changes occur as soon as the voltage changes and circuits operate through the change.

A disadvantage of multiple clock domains is that data generated in one domain and needed in another must cross a domain boundary, potentially incurring synchronization costs. In order to accurately model these costs, we account for the fact that the clocks driving each domain are independent by modeling independent jitter, the variation in the clock, on a cycle-by-cycle basis. Our model assumes a normal distribution of jitter with a mean of zero. In addition, we have analyzed the architectural characteristics of MCD processors which influence overall performance degradation and determined that the dominant characteristic is the ability of the microarchitecture to tolerate pipeline latencies. The accuracy of the domain clock PLLs and the characteristics of the independent clock jitter play a negligible role in overall performance. In fact, both superscalar features (which allow instructions to cross domains in groups) and out-of-order execution (which allows synchronization latencies to be hidden) greatly override the synchronization cost relegating it to an average of less than 3%. The MCD design also has the benefit of reduced clock skew, power, and metal due to the lack of a global skew requirement. However, we assume no such advantages in the analysis that follows.

For the current study we add adaptive structures to the MCD domains. The resulting *adaptive MCD architecture* has a base configuration with small and simple structures running at a very high clock rate. For applications that perform better with additional resources, key structures can be upsized with a corresponding reduction in the clock rate of their domain. Unaffected domains still run at their base high clock rate.

Having adaptable structures and a variable clock means that structures may be safely oversized. The greater capacity (and lower domain frequency) is used only if an application attains a net benefit. Applications that do not require the extra capacity configure to a smaller size and run at a higher frequency. This approach permits the tradeoff between per-domain clock rate and complexity to be made for each application or application phase.

In the front end, the instruction cache and branch predictor are jointly resizable (*i.e.*, each cache configuration is paired with a branch predictor sized to operate at the frequency of the cache). This permits applications with larger instruction footprints, or those that require more branch prediction resources, to be accommodated, albeit at the cost of a lower domain frequency. Similarly, in the load / store domain, the data cache and second level cache are resizable, also as matched pairs. In the integer and floating-point domains, issue queues can similarly be resized to match available ILP. Additional structures could also conceivably be resized; we leave these options for future work.

There are costs associated with supporting adaptive structures. With the smallest sizings, per-domain pipe stage delays are balanced in order to attain the highest clock rate. When the clock frequency is lowered to accommodate the additional delay of an upsized structure, the resulting stage delay imbalance results in a design that is over-pipelined with respect to the particular frequency. The cost is a longer branch mispredict penalty. In our study, the adaptive MCD incurs two additional integer cycles and one additional front-end cycle for branch mispredictions. In addition, the base MCD configuration must have its structures designed for maximum performance in order to achieve the highest possible clock rate. But to support resizing, the smallest struc-

Table 1: L1 data and L2 cache configurations.

L1-D		sub-banks		L2		sub-banks	
size	assoc	adapt	optimal	size	assoc	adapt	optimal
32 KB	1	32	32	256 KB	1	8	8
64 KB	2	32	8	512 KB	2	8	4
128 KB	4	32	16	1 MB	4	8	4
256 KB	8	32	4	2 MB	8	8	4

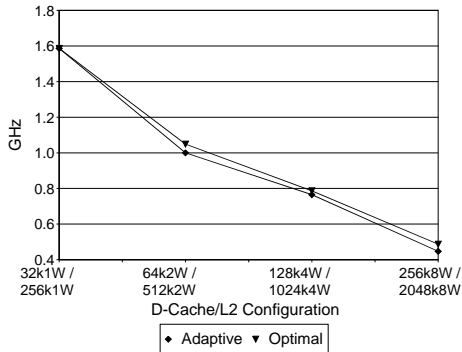


Figure 1: D-cache / L2 frequency versus configuration.

ture size must be a substructure of the larger sizings. Thus, structures may be suboptimal in their large configurations relative to the same size structure that has been optimized for a non-adaptable design.

2.1 Load/store domain

In the load/store domain, the L1 data and L2 caches are 8-way associative, and resized by ways [9]. The base configuration (smallest size and highest clock rate) is a 32 KB direct-mapped data cache and a 256 KB direct-mapped L2 cache. The two caches are upsized together by increasing their associativity. The configurations we consider in our experiments are shown in Table 1.

We use the CACTI modeling tool [29] (version 3.1) to obtain timings for all plausible cache configurations at a given size. The *optimal* columns in Table 1 describe the configurations that provide the fastest cycle time for the given capacity and associativity, without the ability to resize. The *adapt* columns were chosen by adopting the fastest configuration of the minimal-size structure and then replicating this configuration at higher levels of associativity to obtain the larger configurations. This strategy ensures the fastest clock frequency at the minimum configuration, but may not produce the fastest configuration when structures are upsized. Since CACTI configures a 32KB direct-mapped cache as 32 sub-banks, each additional way in the adaptive L1-D cache is an identical 32 KB RAM. The reconfigurable L2, similarly, has 8 sub-banks per 256 KB way. In contrast, the number of sub-banks in an optimal fixed L1 varies with total capacity, and the optimal L2 structure has 4 sub-banks per way for all sizes larger than the minimum.

Frequencies for the various cache configurations, optimal and adaptive, are plotted in Figure 1. The difference between the optimized and adaptive configurations is small: approximately 5%.

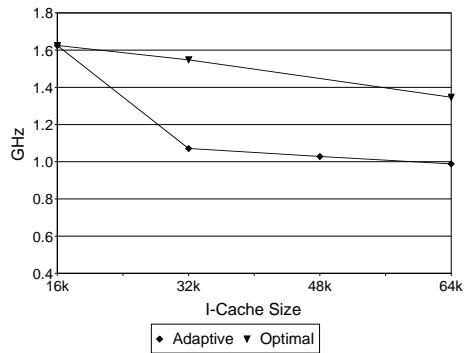


Figure 2: I-Cache Frequency versus Configuration.

2.2 Front end domain

In the front end both the instruction cache and the branch predictor are adaptive. Like the L1-D and L2 caches, they are always resized together. The configurations we consider are shown in Table 2. Like the data caches, the instruction cache adapts by ways, but with associativities of 1, 2, 3, and 4. The branch predictor is a hybrid design with a global shared history (*gshare*) component, a local history component, and a meta predictor to select which component to use. The *gshare* component is a global branch history table (BHT) of 2^{h_g} two-bit counters indexed by the h_g -bit global history. The local history component consists of a pattern history table (PHT) holding the histories for different branches. The table is indexed by the branch PC and returns a h_l -bit wide local history. The h_l -bit wide local history is used to index a local BHT of 2^{h_l} two-bit counters.

As a baseline for comparison, we explored 18 different instruction cache configurations, ranging in size from 4 to 64KB, and in associativity from 1 to 4, to find the best fully synchronous option. Averaged across our suite of 32 MediaBench, Olden, and SPEC2000 applications, a direct-mapped 64KB configuration provides the best performance.

Figure 2 compares the operating frequency of an optimally configured direct-mapped cache to that of our adaptive configurations at various total cache sizes. As is clear on the adaptive curve, there is a large difference in frequency between direct-mapped and 2-way set associative configurations: as much as 31%. Since instruction streams tend to need little associativity, the optimal configuration of a large instruction is direct-mapped and aggressively clocked. This is a clear advantage to the fully synchronous processor that the adaptive design must overcome.

2.3 Integer and floating point domains

In both the integer and floating point domains, the issue queues are resizable from 16 to 64 entries in four increments. The issue queue timings are derived as described by Palacharla *et al.* [21], using the same technology file used by CACTI. Based on the results of Buyuktosunoglu *et al.* [5], we assume that a resizable issue queue suffers no access penalty over a non-resizable issue queue of the same size. A plot of queue frequencies is shown in Figure 3. Note that due to the \log_4 structure of the selection logic and the fact that the selection delay is much larger than the wake-up delay, we suffer a significant frequency decrease when moving from an issue queue with 16 entries (which has 2 levels of

Table 2: Adaptive instruction cache / branch predictor configurations.

I-cache, dynamic			Branch predictor					
Size	Assoc	Sub-banks	h_g	$gshare$ PHT	Meta-predictor	h_l	local PHT	local BHT
16 KB	1	32	14 bits	16384	16384	11 bits	2048	1024
32 KB	2	32	15 bits	32768	32768	12 bits	4096	1024
48 KB	3	32	15 bits	32768	32768	12 bits	4096	1024
64 KB	4	32	16 bits	65536	65536	13 bits	8192	1024

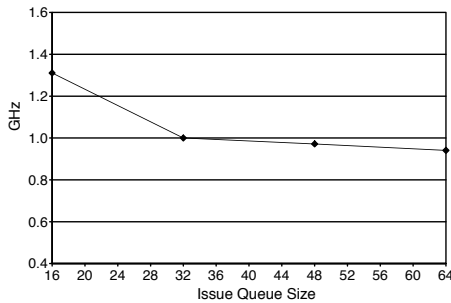


Figure 3: Issue queue frequency analysis.

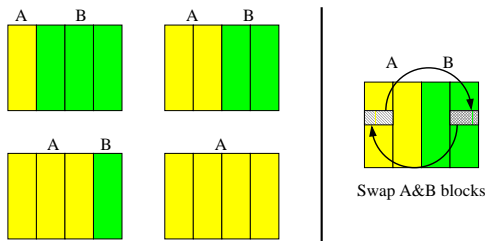


Figure 4: Partitioning options for a 4-way Accounting Cache.

logic) to any larger issue queue up to 64 entries (all of which have 3 levels of logic). As will be explained more fully in Section 5, this frequency effect causes a 16-entry issue queue to be best for most applications.

3. ADAPTIVE CONTROL ALGORITHMS

Both caches and issue queues require some sort of control algorithm to support on-line adaption to program phases. We describe our algorithms in this section. Note that they are not used in the whole-program experiments: for those we pick the configurations that show the best overall behavior based on exhaustive (off-line) exploration of the state space.

3.1 Phase adaptive caches

For our reconfigurable caches we employ the *Accounting Cache* [9] described in our earlier work on energy efficiency. As described in Section 2, this cache adapts by ways. Even when running in a limited number of ways, however, it collects statistics in the remaining ways, allowing us to calculate the number of hits and misses that would have occurred over a given span of time for any of the possible configurations.

The four possible configurations of a 4-way set associative Accounting Cache are shown in Figure 4. In this example, the A partition can be 1, 2, 3, or 4 ways. The B partition is the remaining portion. The A partition is accessed first. If the data block is found it is returned. Otherwise a second access is made to the B partition, which causes blocks to be swapped. Note that the B partition is not considered a lower level of cache. All three caches of the adaptive MCD machine (L1I, L1D, and combined L2) have their own A and B partitions. When we simulate a fully synchronous processor, however, or when we choose a single adaptive configuration for the entire program run, we use only the A partitions: a miss in A skips B and goes directly to the next lower level of the memory hierarchy.

A cache with a small A partition runs at a higher frequency than one with a larger A partition. Ideally one would make the A partition as small as possible without generating a significant number of B accesses. The B partition access latency is an integral number of cycles at the clock rate dictated by the size of the A partition.

As described in detail in previous work [9], we maintain full most-recently-used (MRU) state on cache lines. Simple counts of the number of blocks accessed in each MRU state are sufficient to reconstruct the precise number of hits and misses to the A and B partitions for all possible cache configurations, regardless of the current configuration. Our control algorithm resets the counts at the end of every 15K instruction interval, choosing a configuration for the next interval that would have minimized total access cost in the interval just ended. The hardware components required to perform this calculation are relatively modest. Based on dedicated arithmetic circuits and binary addition trees, we estimate a total of well under 5,000 gates.

3.2 Adaptive issue queues

A special property of the adaptive *Accounting Cache* design is that it avoids exploration of the configuration space when determining the best configuration. This property is also desirable for issue queue control. We introduce a new, deterministic algorithm to measure the inherent ILP of the currently running application, independent of microarchitectural features. The queue control algorithm uses this figure to choose among the four possible queue sizes the one that maximizes effective ILP, normalized to frequency.

A key observation is that the amount of inherent parallelism in the instruction stream can be calculated by immediate dependences of output registers on input registers. The earliest a result can be ready is the latest time of any of its input operands plus the latency of the operation. For this discussion let us assume that all ALU operations have a one cycle latency. As instructions are fetched, their source registers are renamed via the register rename mapping table. With tracking hardware initially reset, all input operands

for the first instruction will have timestamps of zero and the destination register will receive a timestamp of '1'. If the next instruction uses that destination register as an input then its own destination register will be updated with a timestamp of '2', and so on. The maximum timestamp M is continuously recorded during this process.

Tracking continues until N instructions have been fetched, where N represents the queue size of 16, 32, 48, or 64. At that time an estimate of the application's ILP is $\frac{N}{M_N}$. Of course, the division is not actually performed because the numerator is a fixed quantity and the denominators can be compared directly in integer form. When all four estimates have been computed, the control algorithm scales them by the corresponding frequencies and compares them to determine which queue size would have led, in the very recent past, to the highest effective ILP. Dedicated hardware optimized for this task is modest.

The majority of the hardware is in extra storage to hold the timestamps: four bits per register to track the ILP for the 16 entry queue (ILP_{16}), five bits for ILP_{32} , and six bits each for ILP_{48} and ILP_{64} . The tracking intermingles integer and floating point operations, keeping a count of each. A tracking interval ends when either count, N_{INT} or N_{FP} , reaches N . This operation correctly stifles consideration of larger queue sizes that can never be filled for the less dominant instruction type because of resource limitations for the dominant instruction type. With 32 logical integer and 32 logical floating point registers, ILP tracking requires an additional 256 bits for ILP_{16} and 384 bits for ILP_{64} . In our experiments we track the ILP for all queues depths simultaneously, and consider the possibility of resizing as soon as all four counts are available. Alternatively, one could use one set of hardware counters and calculate the four ILP values serially by cycling through the values of N . Counters are reset at the end of each tracking interval.

4. METHODOLOGY

Our simulation environment is based on the SimpleScalar toolset [4] with MCD processor extensions. These extensions include modifications to model an aggressive superscalar processor, *e.g.*, the Register Update Unit (RUU) has been split into separate reorder buffer (ROB), issue queue, and physical register file structures. They also include a heavy re-write of the time management code to emulate separate clocks for each domain, complete with jitter, and to account for synchronization delays on all cross-domain communication. Table 3 contains a summary of the simulation parameters. These have been chosen, in general, to match the characteristics of the Alpha 21264. Much more detail on the MCD extensions can be found in prior papers [25, 26].

We employ a total of 32 benchmark applications: `adpcm`, `epic`, `jpeg`, `g721`, `gsm`, `ghostscript`, `mesa`, and `mpeg2` from the MediaBench suite; `bh`, `bisort`, `em3d`, `health`, `mst`, `perimeter`, `power`, `treeadd`, and `tsp` from the Olden suite; `bzip2`, `crafty`, `eon`, `gcc`, `gzip`, `parser`, `twolf`, `vortex`, and `vpr`, from SPECint2000; and `apsi`, `art`, `equake`, `galgel`, `mesa`, and `wupwise` from SPECfp2000. We run the full application for `bisort`, `treeadd`, and all of the MediaBench applications other than `ghostscript`; the others use instruction windows ranging from 47M-200M instructions. Standard reference inputs were used wherever possible.

The evaluation of the adaptive MCD processor is performed by comparing the relative performance (program

Table 3: Architectural parameters for simulated processor.

Fetch queue: 16 entries
Branch mispredict penalty: 9 front-end + 7 integer cycles (10 + 9 for adaptive MCD)
Decode, issue, and retire widths: 8, 6, and 11 instructions
L1 cache latency (I and D): 2/8, 2/5, 2/2, or 2/- cycles, for A and (optionally) B partitions
L2 cache latency: 12/43, 12/27, 12/12, or 12/- cycles
Memory latency: 80 ns (1st access), 2 ns (subsequent)
Integer ALUs: 4 + 1 mult/div unit
FP ALUs: 4 + 1 mult/div/sqrt unit
Load/store queue: 64 entries
Physical register file: 96 integer, 96 FP
Reorder buffer: 256 entries

execution time) of the best configuration for each application against the fully synchronous processor that provides the best overall performance for our application suite. To find this "best overall" machine we explored a very wide design space: the cross-product of 4 integer issue queue sizes, 4 floating-point issue queue sizes, 4 data/L2 cache organizations, and 16 instruction cache/branch predictor organizations—a total of 1,024 options. (With 32 applications, this portion of our work alone consumed 160 CPU months of simulation time.) The sweep of i-cache configurations in particular ranged from 4KB to 64KB and 1, 2, 3, and 4-way associativity, including many (higher frequency) options not available in the adaptive MCD architecture. The overall best fully synchronous configuration for this suite of benchmarks has a 16-entry integer issue queue, a 16-entry floating-point issue queue, a 64KB direct-mapped instruction cache with a 16-bit *gshare* branch predictor, and a 32KB direct-mapped L1 data cache / 256KB direct-mapped L2 cache. The 64 KB direct-mapped instruction cache is 27% faster than the same capacity in the adaptive MCD design, and the misprediction penalty is substantially lower. The data cache is the smallest and fastest configuration.

5. RESULTS

Shown in Figure 5 is the relative improvement in run time of the Program-Adaptive and Phase-Adaptive MCD processors over the best-overall fully synchronous processor. Program-Adaptive configurations are chosen by per-application exhaustive testing across all possible adaptive MCD configurations. Phase-Adaptive results employ the control algorithms described in Section 3. Many applications achieve a significant performance improvement with the adaptive MCD processor: for `gcc` the Program-Adaptive and Phase-Adaptive processors outperform the fully synchronous processor by 42% and 45%, respectively. For `em3d` the corresponding numbers are 45% and 49%.

When running in Program-Adaptive mode, 85% of our applications use the smallest available integer issue queue. 73% use the smallest floating-point issue queue. 50% use the smallest pair of data caches. 55% use the smallest I-cache. At the same time, there are applications that benefit from larger / slower configurations. Often one structure dominates performance. `Gsm encode` and `decode`, for example, have similar performance for all configurations with a 64KB 4-Way instruction cache. While the very best results are achieved in conjunction with the smallest / fastest in-

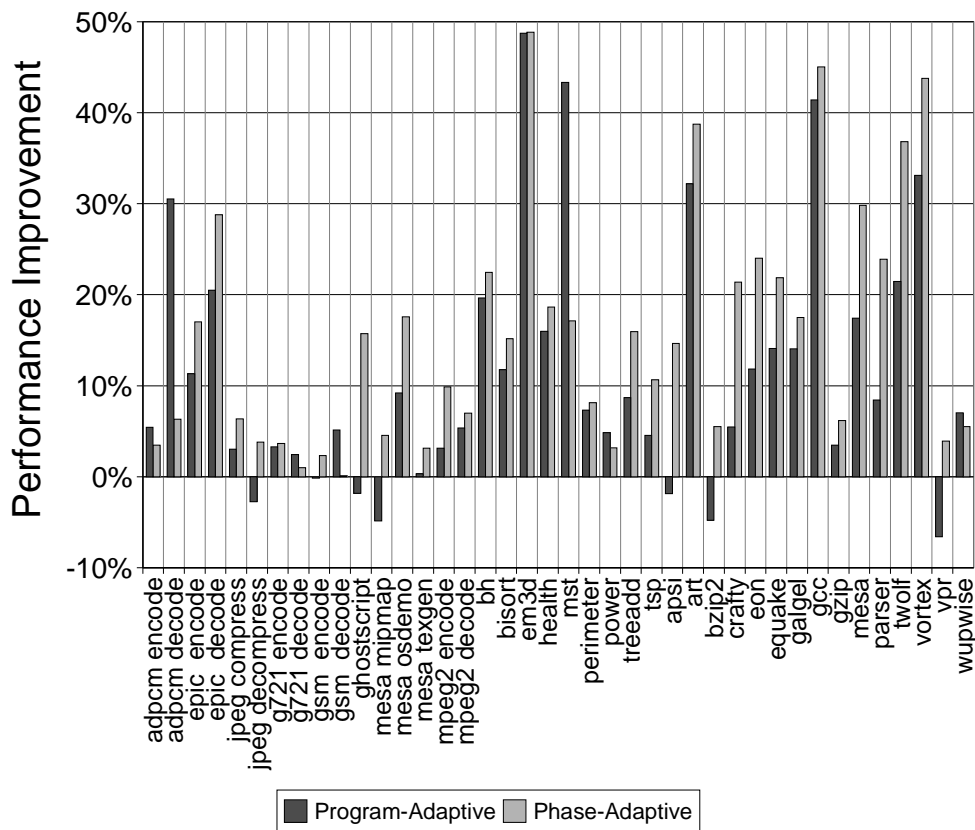


Figure 5: Performance improvement of Program- and Phase-adaptive MCD over fully synchronous.

teger and load/store structures, the differences are minor. Ghostscript, similarly, performs well whenever the instruction cache is larger than 32KB; additional increases yield only marginal improvements.

The bias toward the smallest configuration is due to a number of applications having small computation kernels, high instruction level parallelism, and small data sets; *e.g.*, `adpcm encode`, `adpcm decode`, `bzip2`, and `mpeg2 encode`. For these applications the smaller configuration supplies sufficient capacity for efficient processing; thus, they opt for this configuration due to its higher clock frequency. On the other hand, memory intensive applications such as `em3d` perform best with larger/slower structures that significantly reduce the number of cache misses and avoid the associated long latencies.

There are also applications for which Program-Adaptive MCD performs worse than the fully synchronous processor (`jpeg decompress` 2.7%, `gsm encode` 0.1%, `ghostscript` 1.8%, `mesa mipmap` 4.9%, `apsi` 1.9%, `bzip2` 4.8% and `vpr` 6.6%). The most common source of trouble is an inability to deliver instructions to the execution units fast enough. This in turn stems from the inability to upsize the I-cache without increasing its associativity, which exacts a significant cost in clock rate. For applications that need the larger capacity (but not the increased associativity) there is no way to gain back the performance lost due to the decrease in frequency. For all other applications, the Program-Adaptive MCD improves performance relative to the fully synchronous processor, and in many cases it is a significant improvement

(`adpcm decode` 30.5%, `em3d` 48.7%, `mst` 43.3%, `art` 32.2%, `gcc` 41.4% and `vortex` 33.1%). Overall, a 17.6% performance improvement is achieved.

5.1 Program- vs Phase-Adaptive

In general, Phase-Adaptive outperforms Program-Adaptive. There are no applications in our suite for which the performance of the Phase-Adaptive MCD is lower than that of the fully synchronous processor, and the overall performance improvement increases to 20.4%. In several applications the improvement is significant: `apsi`, `epic encode`, `ghostscript`, `crafty`, `eon`, `mesa`, `parser`, `twolf`, and `vortex`.

A special discussion of the instruction cache is warranted here. All instruction cache configurations are useful for improving performance in the Program-Adaptive results, with nearly half (45%) of the applications configured with an instruction cache other than the smallest. In the Phase-Adaptive case, by contrast, only the smallest instruction cache configuration is used, yet the overall performance is *better* than that of Program-Adaptive. It turns out that adapting at phase changes improves throughput enough to cause the processor to speculate through more branches. This results in a larger number of mispredictions (though not necessarily a higher misprediction *rate*) and more overhead due to pipeline flushes. A smaller, faster instruction cache reduces refill overhead by responding faster, thereby increasing overall performance.

Faster pipeline refill is also why the best instruction cache for the fully synchronous processor is one that is both fast

and large, at 64 KB direct-mapped. The frequency degradation in the adaptive MCD design to go to a larger instruction cache which is set-associative outweighs the performance benefits for these applications. This suggests, as a topic for future work, that an adaptive instruction cache design that resizes by sets rather than ways [23] might be able to gain additional performance benefits.

There are a few notable cases where Phase-Adaptive does not perform as well as Program-Adaptive. One such application is `adpcm_decode`. The cause is the code kernel in the `adpcm_decoder()` function, which contains a series of data-dependent branches that are difficult to predict. The dynamic cache controller speeds up processing by configuring the data caches to their minimal and fastest sizing. The configuration is the proper configuration relative to the cache domain behavior. This results in a high rate of processing and causes more branches to be speculated, often incorrectly. By running with a larger cache the best synchronous design processes instructions more slowly and suffers fewer mispredictions, resulting in a dramatic decrease in pipeline flushes and significantly better performance. This advantage over the Phase-Adaptive case disappears if the problematic branches are replaced with predicated instructions.

Phase-Adaptive mode is less effective with the application `mst` because of periodic short bursts of cache conflicts in the *A* partition. The cache controller responds by increasing the associativity from direct-mapped to 2-way in order to avoid the costly *B* partition accesses. However, the change occurs in the next interval, after the burst, and the cache configuration ends up flipping back to direct-mapped. This pattern repeats, resulting in a slowdown compared to Program-Adaptive mode.

In summary, the adaptive MCD approach is able to exploit the needs of individual applications better than a globally constrained, fully synchronous processor. The Program-Adaptive MCD processor achieves a performance improvement of 17.6% over the best performing fully synchronous processor. Allowing the MCD processor to adapt to application phases further improves performance to 20.4% over the fully synchronous processor and, more significantly, Phase-Adaptive achieves this improvement without the application profiling of Program-Adaptive mode.

6. RELATED WORK

Several manufacturers, notably Intel [16] and Transmeta [12], have processors capable of global frequency and voltage scaling. Global control works well for saving energy in applications with real-time constraints for which the processor as a whole is over-designed [14, 18]. The goal is to save energy with minimum performance *loss*. In [1], Albonesi proposes a *complexity adaptive processor* that adjusts structure capacity while varying global clock frequency and/or access latencies (cycles) for performance optimization.

Childers *et al.* [6] propose trading IPC for clock frequency to save energy. The underlying assumption is that lowering the clock frequency will degrade performance. The user is allowed to select the level of performance degradation to be tolerated. This assumption is correct with a global clock, but in a GALS architecture with decoupled clock domains, decreasing the frequency can give higher performance if resources are scaled up accordingly. Multiple clock domain architectures [3, 15, 25] extend the work of Childers *et al.* by permitting the frequency of each domain to be set indepen-

dently of the others. Semeraro *et al.* [25] adjust frequencies automatically at run time to reduce energy in domains that are not on the critical path of the instruction stream.

Powell *et al.* [23] describe a variable latency data cache. By predicting which way in a set associative cache holds the requested data, an access can be as fast as in a direct-mapped cache. Misprediction forces an additional access. Balasubramonian *et al.* [2] describe a reconfigurable data cache hierarchy whose access time adjusts with the configuration. They assume a globally synchronous processor. We extend this work by decoupling the clock domains and adjusting the data cache hierarchy, instruction cache, branch predictor, and integer and floating-point issue queues.

Reconfigurable issue queues [5, 9, 11, 22] have also been used to reduce energy. Results indicate that applications vary greatly in their resource requirements. The work of Dropsho *et al.* [9] includes the instruction, data, and L2 caches; issue queues; ROB; and register files. While the focus is energy efficiency, the results demonstrate that application resource requirements vary across these structures. Detailed timing of issue queues as it relates to queue size is explored Buyuktosunoglu *et al.* [5] and Palacharla *et al.* [21].

Sasanka *et al.* [24] explore the combination of hardware adaptation (of issue queue size and issue width) and global-chip dynamic voltage scaling for multimedia applications. The goal is to meet frame rate processing requirements while minimizing energy dissipation. Our approach, by contrast, improves *performance* through adaptation and *fine-grain* dynamic *frequency* scaling. We argue that to achieve appreciable speedups, fine-grain dynamic frequency scaling using a GALS approach is necessary to limit the frequency effect of upsizing a structure to the *local unit* level. As we demonstrated in our early work [1], dynamically trading off frequency for complexity in a conventional, single clock, design benefits only those applications with a severe performance bottleneck. This paper demonstrates how the adaptive MCD microarchitecture yields significant speedups across a broad range of applications.

7. CONCLUSIONS

General-purpose processors are designed for good performance across a range of applications. The resulting design point is by necessity a compromise, since applications exercise the microarchitecture resources in widely varying degrees. The *adaptive MCD architecture* offers designers additional dimensions in which the microarchitecture can be optimized to further improve performance. By separating major functionality into separate clock domains, clock rate and complexity tradeoffs can be made independently in each domain. By making the dominant structures in each domain adaptive, these tradeoffs can be made dynamically for each application or application phase.

We demonstrate that an adaptive MCD architecture provides outperforms any fixed, fully synchronous design. When adapting once per application, the average advantage for our suite of 32 standard benchmarks is 17.6%. When adapting automatically in response to program phases, the advantage is 20.4%, despite the branch delay and frequency penalties inherent in adaptability. In future work we plan to explore a wider range of adaptive structures, examine the effects of branch predictor resizing in more detail, and consider resizing the instruction cache by sets instead of (or in addition) to ways.

8. REFERENCES

- [1] D. H. Albonesi. Dynamic IPC/clock rate optimization. In *25th Intl. Symp. on Computer Architecture*, June 1998.
- [2] R. Balasubramonian, D. H. Albonesi, A. Buyuktosunoglu, and S. Dwarkadas. Memory Hierarchy Reconfiguration for Energy and Performance in General-Purpose Processor Architectures. In *33rd Intl. Symp. on Microarchitecture*, Dec. 2000.
- [3] L. Bengtsson and B. Svensson. A Globally Asynchronous, Locally Synchronous SIMD Processor. In *3rd Intl. Conf. on Massively Parallel Computing Systems*, Apr. 1998.
- [4] D. Burger and T. Austin. The SimpleScalar Tool Set, Version 2.0. Technical Report CS-TR-97-1342, U. Wisc.-Madison, June 1997.
- [5] A. Buyuktosunoglu, D. H. Albonesi, S. Schuster, D. Brooks, P. Bose, and P. Cook. A Circuit Level Implementation of an Adaptive Issue Queue for Power-Aware Microprocessors. In *11th Great Lakes Symp. on VLSI*, Mar. 2001.
- [6] B. R. Childers, H. Tang, and R. Melhem. Adapting Processor Supply Voltage to Instruction-Level Parallelism. In *Kool Chips Workshop*, Dec. 2000.
- [7] L. T. Clark. Circuit Design of XScaleTM Microprocessors. In *2001 Symposium on VLSI Circuits, Short Course on Physical Design for Low-Power and High-Performance Microprocessor Circuits*, June 2001.
- [8] A. S. Dhodapkar and J. E. Smith. Managing Multi-Configuration Hardware via Dynamic Working Set Analysis. In *29th Intl. Symp. on Computer Architecture*, May 2002.
- [9] S. Dropsho, A. Buyuktosunoglu, R. Balasubramonian, D. H. Albonesi, S. Dwarkadas, G. Semeraro, G. Magklis, and M. Scott. Integrating Adaptive On-Chip Storage Structures for Reduced Dynamic Power. In *11th Intl. Conf. on Parallel Architectures and Compilation Techniques*, Sept. 2002.
- [10] M. Fleischmann. LongRunTM Power Management. Technical report, Transmeta Corporation, Jan. 2001.
- [11] D. Folegnani and A. Gonzalez. Energy-Efficient Issue Logic. In *28th Intl. Symp. on Computer Architecture*, June 2001.
- [12] T. R. Halfhill. Transmeta breaks x86 low-power barrier. *Microprocessor Report*, 14(2), Feb. 2000.
- [13] A. Hartstein and T. R. Puzak. The Optimum Pipeline Depth for a Microprocessor. In *29th Intl. Symp. on Computer Architecture*, May 2002.
- [14] C.-H. Hsu, U. Kremer, and M. Hsiao. Compiler-Directed Dynamic Frequency and Voltage Scaling. In *Workshop on Power-Aware Computer Systems*, Nov. 2000.
- [15] A. Iyer and D. Marculescu. Power and Performance Evaluation of Globally Asynchronous Locally Synchronous Processors. In *29th Intl. Symp. on Computer Architecture*, May 2002.
- [16] S. Leibson. XScale (StrongArm-2) Muscles In. *Microprocessor Report*, 14(9), Sept. 2000.
- [17] G. Magklis, M. L. Scott, G. Semeraro, D. H. Albonesi, and S. G. Dropsho. Profile-based Dynamic Voltage and Frequency Scaling for a Multiple Clock Domain Microprocessor. In *30th Intl. Symp. on Computer Architecture*, June 2003.
- [18] D. Marculescu. On the Use of Microarchitecture-Driven Dynamic Voltage Scaling. In *Workshop on Complexity-Effective Design*, June 2000.
- [19] V. Milutinovic, D. Fura, and W. Helbig. Pipeline design tradeoffs in a 32-bit Gallium Arsenide microprocessor. *IEEE Trans. on Computers*, 40(11), Nov. 1991.
- [20] J. Muttersbach, T. Villager, H. Kaeslin, N. Felber, and W. Fichtner. Globally-Asynchronous Locally-Synchronous Architectures to Simplify the Design of On-Chip Systems. In *12th IEEE Intl. ASIC/SOC Conf.*, Sept. 1999.
- [21] S. Palacharla, N. Jouppi, and J. Smith. Quantifying the complexity of superscalar processors. Technical Report TR-96-1328, U. Wisc.-Madison, Nov. 1996.
- [22] D. Ponomarev, G. Kucuk, and K. Ghose. Reducing Power Requirements of Instruction Scheduling Through Dynamic Allocation of Multiple Datapath Resources. In *34th Intl. Symp. on Microarchitecture*, Dec. 2001.
- [23] M. Powell, A. Agrawal, T. N. Vijaykumar, B. Falsafi, and K. Roy. Reducing set-associative cache energy via selective direct-mapping and way prediction. In *34th Intl. Symp. on Microarchitecture*, Dec. 2001.
- [24] R. Sasanka, C. Hughes, and S. Adve. Joint Local and Global Hardware Adaptations for Energy. In *10th Intl. Conf. on Architectural Support for Programming Languages and Operating Systems*, Oct. 2002.
- [25] G. Semeraro, D. H. Albonesi, S. G. Dropsho, G. Magklis, S. Dwarkadas, and M. L. Scott. Dynamic Frequency and Voltage Control for a Multiple Clock Domain Microarchitecture. In *35th Intl. Symp. on Microarchitecture*, Nov. 2002.
- [26] G. Semeraro, G. Magklis, R. Balasubramonian, D. H. Albonesi, S. Dwarkadas, and M. L. Scott. Energy-Efficient Processor Design Using Multiple Clock Domains with Dynamic Voltage and Frequency Scaling. In *8th Intl. Symp. on High-Performance Computer Architecture*, Feb. 2002.
- [27] T. Sherwood and B. Calder. Time Varying Behavior of Programs. Technical Report UCSD-CS99-630, U. Cal. San Diego, Aug. 1999.
- [28] E. Sprangle and D. Carmean. Increasing Processor Performance by Implementing Deeper Pipelines. In *29th Intl. Symp. on Computer Architecture*, May 2002.
- [29] S. J. E. Wilton and N. P. Jouppi. CACTI: An Enhanced Cache Access and Cycle Time Model. *IEEE J. of Solid-State Circuits*, May 1996.
- [30] B. Xu and D. Albonesi. Runtime reconfiguration techniques for efficient general purpose computation. *IEEE Design and Test of Computers*, Jan. 2000.