

Reconfiguring Issue Logic for Microprocessor Power/Performance Throttling

Edwin Olson eolson@mit.edu, Andrew Menard armenard@mit.edu

Abstract—A growing need for computational power in mobile devices has spawned increased interest in low-power microprocessors. Some low-power applications require very high performance, such as real-time video decoding on Personal Digital Assistants. A growing body of work has examined how to provide this high performance when needed, while throttling performance so that power consumption can drop to very low levels when performance is not required. Observing that the issue logic in an out-of-order microprocessor consumes a significant amount of power, several groups have attempted to modify this part of the processor so that it can dynamically enter a low-power mode. We have revisited these topics and our work shows that simple approaches to modifying issue logic fail to reduce the average energy per instruction. We also look at the possibility of including a low-power single-issue processor on the same die as a high-performance multiple-issue processor. Swapping between these two processors allows a dynamic tradeoff between power and performance.

Index terms—Issue Window, Issue Logic, Out-of-Order, Low Power, Power/Performance Throttling

I. INTRODUCTION

Much of the thrust of recent computer architecture work has been in search of increased performance. As transistor budgets increased, more and more technologies from mainframes were incorporated in microprocessor designs. The product of this evolution was high performance microprocessors that sacrificed power consumption for performance. With the emergence of low-power markets, these speed demons have been retrofitted to consume less power by incorporating clock gating, voltage scaling, and more recently, dynamic resizing of key architectural features such as the issue window.

Many existing techniques for reducing power are well established and extremely effective, including dynamically reconfiguring the cache[1] and voltage scaling.[8] Reducing the supply voltage of a microprocessor has a roughly linear effect on performance (due to weaker

electric fields) but a squared effect on power dissipation (since power consumption is proportional to $\frac{1}{2} \cdot \text{frequency} \cdot CV^2$).

When comparing two architectures for power efficiency, it is tempting to use a metric such as energy/instruction or the power delay product. However, one must take into account the required performance level. A processor with seemingly poor energy/instruction characteristics that has more performance than required can be run at a lower voltage thus reducing performance and energy/instruction. This processor might then be much more attractive.

While voltage scaling is a very good way of providing additional power/performance modes, it has its limits. When operating voltage approaches the threshold voltage of the transistors, the performance of the transistors begins dropping off much faster than linearly. As threshold voltages are reduced, leakage currents increase, which, in turn, increases power consumption. The Semiconductor Industry Association predicts that in the year 2005, supply voltages for low power applications will be 0.9-1.2V, compared to a typical modern supply of 1.8V [17]. Even while this reduction implies significant power consumption, the SIA predicts a net increase in total power requirements for battery-operated devices of 70%. Clearly there is a need for additional power/performance throttling mechanisms.

It has been observed that one major power drain in modern out-of-order processors is the issue logic; every clock cycle, each instruction in the issue queue must be checked to see if it can be dispatched. Retired instructions broadcast the availability of new operands on long bit lines across the entire issue window. Some processors, such as the Alpha 21264, compact the issue queue in order to implement an oldest-first priority algorithm, and this process requires even more energy. In the 21264, between 18 and 46 percent of the total power of the processor is consumed by the issue logic [6].

Thus, methods of scaling back the size of the issue window and the number of instructions issued each cycle have been proposed in order to minimize this source of power

The authors are graduate students at the Massachusetts Institute of Technology, Cambridge, MA.

consumption at the cost of reduced performance. These methods are compatible with cache disabling and voltage scaling; for maximum reduction in power consumption the power management software could simultaneously reduce the voltage to the lowest possible level, disable parts of the cache, and reduce the issue window size, or it could find intermediate power/performance points by doing only one or two of these optimizations. We also consider an alternate scheme of bypassing complex issue logic completely. We propose to do this by placing an in-order, single-issue core alongside the out-of-order multiple-issue core, with the OS able to swap between them, thus avoiding the issue logic altogether when necessary. In this paper we do not have a required performance level that the chip must achieve. Instead consider the metric of power consumed per instruction per cycle.

Several studies have shown that relatively simple modifications can allow an operating system to do performance throttling without spending an excessive amount of time profiling the code being executed, [11] and that relatively simple hardware structures can also monitor performance needs.[13]

II. METHODOLOGY

In order to conduct our study, we needed to measure the impact of changing architectural resource sizes, such as the number of slots in the issue window, on both power and performance. The SimpleScalar toolset provides detailed performance simulators [4]. SimpleScalar provides a performance simulator using a relatively unique microarchitecture built around a “Register Update Unit”, an architectural resource combining the functions of the issue window and the register renaming unit. This is somewhat unfortunate, since it does not correlate well to actual chips.

However, the results we receive from these studies can still yield insight into the effects of scaling architectural features, and the relative results are still meaningful. Many other architectural studies have also used SimpleScalar, so our results can be directly compared with those. Future work may involve repeating our studies with a model more closely resembling commercially successful architectures.

SimpleScalar does not provide a mechanism for measuring power. However, several research groups have added power models to SimpleScalar, such as David Brooks' Wattch tool [3] and the Cai-Lim models [9]. Wattch's models are better suited to our study because its models are heavily parameterized and are therefore capable of reflecting various changes in configuration without needing to create SPICE models for each variation. We used version 1.02 of the Wattch power.c model.

Power estimation tools like Wattch and the Cai-Lim models have recently been the subject of considerable scrutiny.[10]. Ghiasi, Grunwald and others have shown that not only are direct comparisons of energy measurements hopeless due to disagreement, but even relative comparisons often fail to agree. A key problem rests in the fact that an architectural description simply doesn't contain an adequate amount of information to properly estimate power, and even reasonable parameterized models quickly become unrealistic when the parameters are adjusted beyond a limited range. For example, the Wattch CAM model used for the RUU structure is a reasonable model for a 16 entry structure, but if the structure had 256 entries, it would have been implemented in a completely different way such as being banked. Therefore, we have limited our study to modifying parameters by only small factors.

We consider 4 issue and 8 issue processors with varying RUU sizes. The other characteristics of our processors are listed in Table 1. SimpleScalar allows many parameters to be adjusted, but we only changed a handful. Table 1 is primarily a list of non-default settings.

Table 1.

	4 issue	8 issue
Decode Width	4	8
Commit Width	4	8
Load Store Queue Size	8	8
Integer ALUs	4	6
Integer Multipliers	1	2
FP ALUs	4	4
FP Mul/Div	1	2
Memory Ports	2	4

Our benchmarks are derived from the SpecInt95 suite. Due to the limited speed of the SimpleScalar simulator (about 90k instructions per second), it was impractical to run the entire suite, or even an entire single benchmark. Instead, as is the common practice in the simulator field, reduced input sets were used. These input sets take substantially less time to test, but still exercise the processor in ways similar to the official input sets. Therefore, the performance data we generated cannot be compared to actual SpecInt scores, but we are primarily interested in the relative performances of our various models.

Table 2.

Benchmark	Input
Li	nqueens 6
Perl	test.in
compress95	5000 q 2131

For all of these benchmarks, the kernel of the program, rather than initialization code, dominated the runtime. In addition, the simulator is completely deterministic, so there is no need to repeat simulations and average scores.

III. DETERMINING OPTIMAL RUU CAPACITY

Understanding the optimal size for the Register Update Unit is extremely important when determining its actual capacity. Several factors influence this optimal size. The goal of the RUU is to always have enough instructions ready to feed the available functional units. As the number of functional units increases, the size of the RUU should intuitively increase to provide more candidate instructions. However, due to data dependencies, it is often the case that the number of instructions that can be fetched is greater than the number that can be issued. We want the RUU to hold a certain “surplus” of instructions so that when an instruction miss occurs and fetch rate drops to zero, the functional units can be kept busy, but there’s no reason to make the RUU unreasonably large.

A. Bounds on RUU Usage

Our first experiment’s goal was to determine an absolute upper bound on the size of the RUU. We configured SimpleScalar to use an extremely large RUU and made modifications to SimpleScalar to collect statistics on the size of the RUU every cycle. The resulting structure could hold enough instructions to keep the functional units busy for dozens of cycles, and is therefore excessive. However, it does provide an upper bound on the size of the RUU from which to work from.

Figure 1.




Figure 2.





Figure 1 shows that for all three benchmarks, the RUU almost never contains more than 32 instructions at either issue width. Making an RUU any larger than 32 would serve no function; the entries would be empty almost all the time.

When the RUU’s physical size is bounded, the RUU usage closely mirrors the unlimited case, except the RUU “saturates”. In figure 3, we see that a 16 entry RUU has almost exactly the same occupancy characteristics when occupancy is between 0 and 15. The 16 entry RUU is fully occupied about as often as the unlimited RUU has 16 or more entries. This is as would be expected, and though Figure 3 uses the li benchmark, the other benchmarks have the same behavior.

Figure 3.



B. IPC vs. RUU size

The important question now is: if the RUU capacity is limited beyond the ideal case, what happens to performance? We measured performance in terms of Instructions Per Cycle (IPC), since we cannot accurately

determine changes in clock period from within SimpleScalar.

Figure 4.

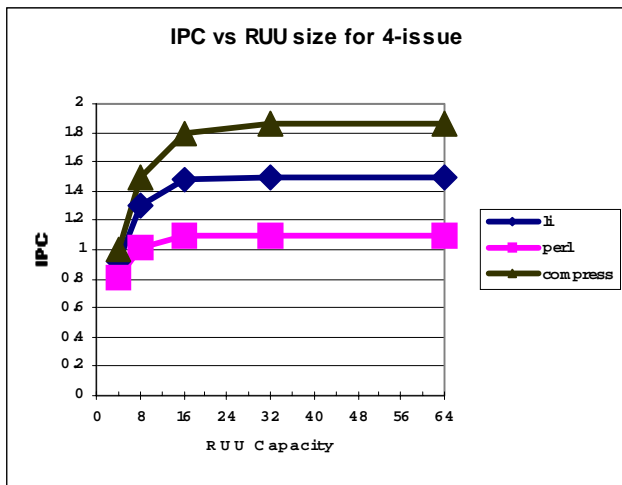


Figure 5.

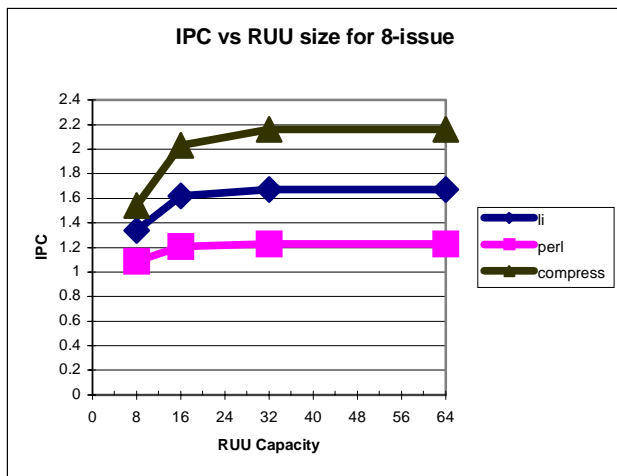


Figure 4 shows the performance of the processor, in terms of IPC, versus the capacity of the RUU. We notice immediately that the performance of the processor for compress and perl is *very* similar for RUU capacities of 16 and 32 for a 4-issue processor. There’s a small increase for li. As we expected, there is almost no benefit in scaling the RUU beyond 32.

If we consider an 8-issue machine, we would expect the performance of the processor to drop off more rapidly than the 4-issue with decreasing RUU capacity. This is because the RUU could be depleted (potentially) twice as quickly, and the processor is therefore more likely to be unable to keep its functional units busy. We see precisely this behavior in Figure 5; there is a noticeable performance difference for both li and perl between RUU capacities of 16 and 32.

Some research groups have proposed dynamically varying the issue window capacity. [14] It is obvious that a parameterized model of an RUU is likely to predict substantially greater power consumption for a 32-entry RUU than a 16-entry RUU. We must resist the temptation to declare this to be an efficient mechanism for throttling power/performance. While performance is affected, a power-conscious architect is unlikely to make the RUU so much larger for such a miniscule return. This is an uninteresting regime since the IPC vs RUU size curve is essentially flat.

C. Relationship between energy and RUU size

However, an interesting question still remains. What happens to energy per instruction statistics as we decrease the RUU well into the region of decreased performance? It might be a good idea to allow a processor to dynamically decrease its RUU size, for example from 16 to 8, if the decrease in power offsets the decrease in performance.

Using the Wattch tool, we measured the power consumption of the processors and calculated the average energy per instruction assuming optimal clock gating (Wattch’s cc3 models).

Table 3.

Structure	4x4	4x8	4x16	4x32	4x64
Energy/Inst (li)	15.8	13.0	11.8	12.8	14.1
Energy/Inst (perl)	16.5	14.3	13.6	14.7	16.1
Energy/inst (compress)	14.4	11.5	10.6	11.3	12.5


Table 3 shows the average energy per instruction for each benchmark, for various RUU capacities of a 4-issue processor. We already expected the 4x32 and 4x64 configurations to be suboptimal, since the RUU is essentially oversized. It’s interesting, however, that the cost of executing instructions actually increases when the RUU is shrunk below 16 entries. While the power consumption of the issue logic is going down with decreasing RUU capacity, the performance is dropping super-linearly.

We can also see that we’re spending more energy per instruction on codes with less inherent parallelism (perl in particular). This makes sense since there are a lot of hardware resources in an out-of-order superscalar processor looking for parallelism to exploit, but there’s simply very little parallelism to be found. This overhead cost is being amortized over very few issued instructions every cycle, and thus the average energy per instruction is higher.

We'll also note that we don't trust the power numbers for the extreme configurations of RUU (x4 and x64) since they comprise a significant factor of deviation from Wattach's baseline capacity.

The breakdown of power consumption is shown in figure 6. The patterns of power consumption are similar for all three benchmarks, so we show only the li case. One component that is consuming conspicuously more power as the RUU size is increased is the RUU itself (denoted as 'window'). Somewhat unexpected are the increases in energy in other areas of the chip. As discussed earlier, when the RUU size is increased, the total number of instructions (committed + speculated) executed increases 12-23%, depending on the benchmark. This causes increased activity in almost all of the major functional blocks. In addition to the window energy, we see significant increases in the clocking energy, the load store queue, and the result bus.

Figure 6.



In Table 4, we have energy per instruction statistics for an 8 issue processor. We see very similar trends as in the 4-issue processor. As with the 4-issue case, we observe that the 16-entry RUU is the minimum energy per instruction point.

Table 4.

Structure	8x8	8x16	8x32	8x64
Energy/Inst (li)	13.8	12.5	13.4	14.9
Energy/Inst (perl)	15.1	14.7	15.8	17.6
Energy/inst (compress)	12.4	11.4	11.9	13.3

IV. OTHER LOW-POWER MODIFICATIONS TO COMPLEX PROCESSORS

It seems as though scaling a processor's issue window will not provide the power/performance throttling we would like. There are many potential functional units that can be targeted for energy reduction, but other difficulties arise.

We see from Figure 6 that the register file consumes a significant percentage of power. SimpleScalar's RUU structure works in its favor for minimizing the complexity of the register file by incorporating the renamed registers within the issue window and maintaining a separate (and smaller) architectural register file, whereas in a mainstream design the register file often contains both the renaming and architectural registers. In the latter case, the register file is both physically larger and may have additional ports, consuming even more power.

V. LOBOTOMIZING AN OUT-OF-ORDER PROCESSOR

Our group considered several mechanisms for dynamically "lobotomizing" an out-of-order processor in order to provide new power/performance points. Our initial approaches mirrored those of other groups—dynamically resizing issue logic, but our study of the effects of RUU sizing discussed previously in this paper made this seem problematic.

Our second idea was to disable most of the logic accompanying the out-of-order issue logic, essentially causing instructions to be issued in-order. We modeled this in SimpleScalar by disabling out of order execution and speculative execution, then reducing the size of the RUU to one. We expected extremely poor performance based on our previous experiments in RUU sizing. We measured an effective IPC of 0.57. The power numbers returned by Wattach are not reported here since their accuracy with the oddly-sized models cannot be relied upon. The major cause of the poor performance is the very high latency of an out-of-order processor (compared to a simple pipelined machine) which causes many stalls when dependent series of instructions are run. An out-of-order machine spends a lot of time in order to find and exploit parallelism, and dramatically reducing the RUU's capacity causes most of this work to be wasted since it eliminates the possibility of having many instructions executing simultaneously.

One idea that our group considered was completely bypassing the complicated issue logic of a complex microprocessor. If the renaming logic and issue window were bypassed completely—if a single instruction was passed immediately from the instruction fetch stage to the register file read stage—the latency of instructions would be substantially shorter and the performance would increase substantially. It would be tempting to use a banked register file as well, so that when register renaming was turned off, access to the register file would all come from a smaller "architectural register" register file. However, unless a completely separate register file was used for the low-power mode of operation, each bank of the register file

would likely have far more ports than would be necessary for a single-issue processor, and this would put a bound on the amount of power savings that could be achieved. It would be worthwhile to build a model of this and simulate it in detail in the future.

VI. USING A COMPLETELY SEPARATE CORE

Since simply scaling back the size of the issue window does not seem to be an obvious win, we also considered eliminating it altogether. A large, complex, out-of-order core could be used when high performance was required, or a small, simple, in-order pipelined core could be used when low-power operation was needed. Our intuition suggested that a simpler core would likely have much lower energy/instruction given the same technology. Also, since the cores are completely separate, each can be optimized separately. If the on-die caches and other circuits could be used for both cores, the overhead in die area for a small in-order core would be very small.

We wished to get the most accurate data possible, so rather than using Wattch and SimpleScalar, we opted to conduct a survey of commercially available processors, looking for processor families where there is both an in-order and an out-of-order implementation of the same ISA in the same technology, and of processors that support some form of voltage scaling, for a comparison.

IBM's PowerPC line includes the model 440 CPU, a dual-issue, 7-stage pipeline machine, and the 405 CPU, a single-issue, 5-stage pipeline machine, both implemented in the same .18 micron copper process [15,16]. The 440, operating at 550MHz, consumes approximately 1.0W of power, and performs at 1000mips on the Dhrystone 2.1 benchmark, while the 405 operating at 266MHz consumes approximately 0.5W of power while performing 375mips on the same benchmark. Thus, the energy used per instruction on the 440 is approximately 1mj, while on the 405 it is approximately 1.3mj. This is a very disappointing result; the faster processor is actually using less energy per instruction, so clearly you would benefit more from an approach like voltage scaling to reduce total energy used across a calculation, or just use the faster processor until the calculation is finished and then put it into a sleep mode, both of which also avoid the significant area overhead of the dual processor approach.

The Intel Pentium III mobile versions utilize voltage scaling from 1.6V to 1.3V to achieve a more than 50% reduction in power consumption while still achieving 70% of the performance. Intel's Xscale line of StrongARM-compatible chips uses voltage scaling at a much finer granularity to go from consuming 450mW at 800MHz to only 40mW at 150MHz[17]. The Transmeta corporation's Crusoe line of chips use dynamic voltage and frequency scaling from 600 MHz at 1.6V to 300MHz at 1.2V, achieving a significantly better than linear drop in power

consumption for a linear drop in speed [18]. Each of these chips manages a significantly better than linear power/performance tradeoff, and most new chips aimed at mobile computing feature a low-power sleep mode, and even the relatively simple throttling mechanism of cycling the processor into and out of this mode achieves a nearly linear power/performance tradeoff.

This demonstrates that in mainstream processors with today's technology, voltage scaling is still the best approach to power/performance throttling.

VII. CONCLUSIONS

While the bulk of recent computer architecture research has focused on increasing performance, many modern applications require both high performance and low power. Techniques such as voltage scaling and clock gating are well established as ways to reduce power consumption without adversely affecting performance. In this paper, we considered two techniques for switching between a high-performance mode and a lower-power, low-performance mode. Dynamically changing the processor's issue width seemed promising initially but yielded poor results when considering the power consumed per instruction. Installing a small in-order core alongside the out-of-order core also fails to yield an overall benefit in terms of power used.

On a high-performance processor such as the Digital Alpha 21264, between 18 and 46 percent of the total power consumed by the processor goes to the issue logic. This suggests bypassing or reducing the issue logic as a route to minimizing power consumption if performance is not a concern. Using Wattch and SimpleScalar, we first determined that there is a maximum size for the SimpleScalar register update unit, or RUU. Increasing the size of this structure, roughly analogous to a real microprocessor's register renaming logic and issue window, yields no performance gains if this maximum size is exceeded. We then examined the power used by the RUU, and found that the performance of the processor drops faster than its power requirements when the RUU size is decreased. Thus, the power per instruction is optimal when the RUU is at this maximum size, about 16 instructions for either a 4- or 8-issue processor.

Because of this result, changing the size of the issue window on an actual microprocessor would not be beneficial in terms of power consumed. Real-time applications require a certain number of instructions to be executed in a particular time frame. Our results show that the minimal power is used by doing this computation using the full power of the microprocessor, and then switching to a very low power sleep mode in which no instructions are executed for the remainder of the time period. Attempting to change the issue width would result in more power

being used per instruction; since in this scenario the number of instructions executed per unit time is constant, this results in more power used per unit time.

Given these discouraging results, we also considered the possibility of including a completely separate in-order core on the die of a larger microprocessor. We hoped that the in-order core would be small enough to fit into a modern superscalar machine without significantly impacting the layout. While the in-order machine would have much lower performance, it ideally would have had a much lower power usage since it lacked the expensive register renaming and instruction reordering logic present on the out-of-order machine.

A survey of real processors suggests that this unfortunately is not the case. In the PowerPC family, the in-order 405 chip consumes 30% more power than the superscalar 440. Intel's microprocessors have also avoided these techniques, relying instead on traditional techniques to obtain a 50% power savings in mobile versions of their CPUs.

In this paper, we have shown that neither modifying a processor's issue width nor adding a separate in-order core offers a possibility for a low-power mode as an alternative to a high-performance mode. We hope to be able to use better simulation tools to examine some of the options presented here, including completely bypassing the issue logic, using a separate register file for in-order execution, and including a separate in-order core on chip. However, current widely-used techniques, such as voltage scaling and clock gating, appear to offer the best power savings currently available for low-power applications.

VIII. REFERENCES

- [1] David H. Albonesi, "Dynamic IPC/Clock Rate Optimization," 25th International Symposium on Computer Architecture, 282--292, June, 1998
- [2] W. Ye and N. Vijaykrishnan and M. Kandemir and M. J. Irwin, "The Design and Use of SimplePower: A Cycle-Accurate Energy Estimation Tool," 37th Design Automation Conference, 340--345, June 2000
- [3] David Brooks and Vivek Tewari and Margaret Martonosi, "Wattch: a framework for architectural-level power analysis and optimizations," 27th Annual International Symposium on Computer Architecture, 83--94, June, 2000
- [4] Doug Berger and Todd M. Austin, "The SimpleScalar Tool Set, Version 2.0," June, 1997},
- [5] David H. Albonesi, "The Inherent Energy Efficiency of Complexity-Adaptive Processors," 1998 Power-Driven Microarchitecture Workshop, held at the 25th International Symposium on Computer Architecture, 107--112}, June, 1998
- [6] Michael K. Gowan and Larry L. Biro and Daniel B. Jackson, "Power considerations in the design of the {Alpha} 21264 microprocessor," "35th Annual Conference on Design Automation", 726--731, June, 1998
- [7] R. Y. Chen and M. J. Irwin, "An Architectural Level Power Simulator," 25th International Symposium on Computer Architecture, June, 1998
- [8] T. Pering and T. Burd and R. Broderon, "Dynamic Voltage Scaling and the Design of a Low-Power Microprocessor System," 1998 Power-Driven Microarchitecture Workshop, held at the 25th International Symposium on Computer Architecture, 107--112, June, 1998
- [9] G. Cai and C. H. Lim, "Architectural level power/performance optimization and dynamic power estimation," MICRO32, November, 1999
- [10] Soraya Ghiasi and Dirk Grunwald, "A Comparison of Two Architectural Power Models," Ninth International Conference on Architectural Support for Programming Languages and Operating Systems, November, 2000
- [11] L. Benini and A. Bogliolo and S. Cavallucci and B. Ricco, "Monitoring System Activity for OS-Directed Dynamic Power Management," International Symposium on Low Power Electronics and Design, August, 1998
- [12] Vivek Tiwari and Deo Singh and Suresh Rajgopal and Gaurav Mehta and Rakes Patel and Franklin Baez, "Reducing Power in High-performance Microprocessors," 35th Annual Conference on Design Automation, June, 1998
- [13] Roberto Maro, Yu Bai, and R. Iris Bahar, "Dynamically Reconfiguring Processor Resources to Reduce Power Consumption in High-Performance Processors"
- [14] Alper Buyutosunoglu, Stanley Schuster, David Brooks, Pradip Bose, Peter Cook, and David Albonesi, "An Adaptive Issue Queue for Reduced Power at High Performance".
- [15] IBM Product Datasheet for the PowerPC 440 Core
- [16] IBM Product Datasheet for the PowerPC 405 Core
- [17] Silicon Industry Association, "International Technology Roadmap for Semiconductors, 1999 Edition", http://public.irirs.net/files/1999_SIA_Roadmap/Home.htm
- [18] Intel Xscale Microarchitecture Technical Summary <http://developer.intel.com/design/intelxscale/XScaleDatasheet4.htm>
- [19] Marc Fleishmann, "Crusoe Power Management," HotChips 12