

Architectural Tradeoffs for Low Power

Vojin G. Oklobdzija*

Integration
Berkeley, CA 94708
<http://www.integr.com>

*Electrical Engineering Department
University of California, Davis, CA 95616
<http://www.ece.ucdavis.edu/acsel>

Abstract

A study of several RISC, DSP and embedded processors was conducted. It has been shown that the transistor utilization drops substantially as the complexity of the architecture increases. Simple architecture that are enabling full utilization of technology are favorable as far as energy efficient design styles are concerned. The results favor simple architectures that leverage performance improvements through technology improvements.

1. Introduction

Demand for reducing power in digital systems has not limited to systems which are required to operate under conditions where battery life is an issue. The growth of high-performance microprocessors has also been constrained by the power-dissipation capabilities of the package using inexpensive air-cooling techniques. That limit is currently in the neighborhood of fifty watts. However, the increasing demand for performance (which has been roughly doubling every two years) is leaving an imbalance in the power dissipation increase, which is growing approximately at 10 Watts per year.

This growth is threatening to slow the performance growth of future microprocessors. The “*MOS ULSIs are facing a power dissipation crisis*” in the words of Kuroda and Sakurai [1]. The increase in power consumption for three generations of Digital Equipment Corporation, “Alpha” architecture high-performance processors is given in Fig. 1.

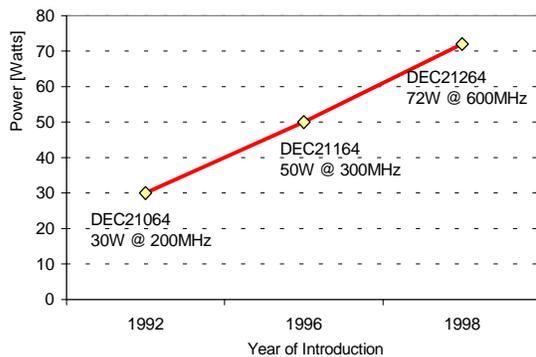


Fig. 1. Power increase for three generations of DEC ‘Alpha’ processor

2. Comparative Analysis

Most of the improvement on power savings is gained by technology. Scaling of the device and process features, lowering of the threshold and supply voltage result in an order of magnitude savings in power. Indeed, this resulting power reduction has been a salient achievement during the entire course of processor and digital systems development. Had this not been the case then the increase in power from one generation to another would have been much larger limiting the performance growth of microprocessors much earlier.

The technology amounts for approximately 30% improvement in gate delay per generation. The resulting switching energy CV^2 has been improving at the rate of 0.5 times per generation. Given that the frequency of operation has been doubling for each new generation the power factor $P = CV^2f$ remained constant ($0.5 \times 2 = 1.0$). It is the increase in complexity of the VLSI circuits that goes largely uncompensated as far as power is concerned. However, it is estimated that the number of transistor has been tripling for every generation. Therefore, the expected processor performance increase is 6 times per generation (two times due to the doubling of processor frequency multiplied by the three times increase in the number of transistors).

The fact that the performance has been increasing four times per generation instead of six is a strong indication that the transistors are not efficiently used. What that means is that the added architectural features are at the point of diminishing returns.

This diminishing trend is illustrated in Table 1 which compares a transition from a dual-issue machine to a 4-way-super-scalar for the IBM PowerPC architecture.

All three implementations of the PowerPC architecture are compared at the same frequency of 100MHz. The performance of PowerPC 620, as well as power consumption has been normalized to what it would have been at 100MHz. We can observe that the power has more than doubled and quadrupled respectively in transition from a relatively simple implementation (601+) into a true super-scalar 620. The respective performance has also improved by 50 to 60% (integer) and 30 to 80% (floating-point). However, the number of Specs/Watt has gone down dramatically-- one and two times as compared to 601+. Given that all the data for all the three implementation has been compared at 100MHz, we

Table 1. Comparison of PowerPC performance / power transition[8,9]

Feature	601+	604	620	Diff.
Frequency MHz	100	100	133 (100)	same
CMOS Process	.5u 5-metal	.5u 4-metal	.5u 4-metal	~same
Cache Total	32KB Cache	16K+16K Cache	64K	~same
Load/Store Unit	No	Yes	Yes	
Dual Integer Unit	No	Yes	Yes	
Register Renaming	No	Yes	Yes	
Peak Issue	2 + Br	4 Insts	4 Insts	~double
Transistors	2.8 Million	3.6 Million	6.9 Million	+30% /+146%
SPECint92	105	160	225 (169)	+50% /+61%
SPECfp02	125	165	300 (225)	+30% /+80%
Power	4W	13W	30W (22.5W)	+225%/+463%
Spec/Watt	26.5/31.2	12.3/12.7	7.5/10	-115%/-252%

are indeed comparing the inverse of Energy-Delay product which is a true measure for power efficiency of an implementation as shown in [7].

The comparable inefficiency in power-performance factor in transition from single-issue to a super-scalar for MIPS processor architecture is shown in Table 2. The comparison shows a 31% decrease in power efficiency for the integer code but a 23% improvement for the floating-point.

Table 3 shows that the best trade-off between performance and power has been achieved in DEC Alpha 21164 implementation of their "Alpha" architecture. The table shows comparable efficiency for MIPS, PowerPC and HP processor implementations, slightly better for Sun UltraSPARC and substantially better power efficiency for Digital 21164.

The power efficiency of DEC 21216 was achieved through very careful circuit design, thus eliminating much of the inefficiency at the logic level. This was necessary in order to be able to operate at the frequency that is twice as high compared to other RISC implementations. However, no architectural features, other than their very careful implementations are contributors to the power efficiency of DEC 21164.

It is interesting to compare what a particular improvement means in terms of power. In Table 4 we are comparing the effect of increasing the cache size for IBM 401 and 403 processors. The measurement is normalized to 50MHz. The power-efficiency has dropped by a factor of close to two, resulting from increasing the caches. Similar findings are confirmed in the case of PowerPC architecture where the decrease in power efficiency is 60% as shown in Table 5.

Table 2. Transition from single issue MIPS R5000 to MIPS R10000 implementation of MIPS architecture[8,9]

Feature	MIPS R10000	MIPS R5000	Diff.
Frequency	200MHz	180MHz	~same
CMOS Process	0.35 /4M	0.35 /3M	
Cache Total	32K/32KB Cache	32K/32K Cache	~same
Load/Store Unit	Yes	No	
Register Renaming	Yes		
Peak Issue	4 Issue	1+FP	
Transistors	5.9 Million	3.6 Million	+64%
SPECint95	10.7	4.7	+128%
SPECfp95	17.4	4.7	+270%
Power	30W	10W	200%
SPEC/Watt	0.36/0.58	0.47/0.47	-31%/23%

Metrics:

Horowitz et al.[7] introduces Energy-Delay product as a metric for evaluating power efficiency of a design.

An appropriate scaling of the supply voltage results in a lower power, however, at the expense of the speed of the circuit. The energy-delay curve shows an optimal operation point in terms of energy efficiency of a design. This point is reached by various techniques, which are all being discussed in this paper.

The fabrication technology seems more important for the energy-delay that the architectural features of the machine. This finding is consistent with the fact that the processors' performance has been increasing four-fold per generation. Though we would expect a six-fold increase in performance: the frequency has been doubling per generation and the number of transistor tripling. This shows that the transistors have not been used efficiently and that the architectural features that are consuming this transistor increase have not been bringing a desired effect in terms of the energy-efficiency of the processors.

Power Tradeoffs in DSP and Embedded Systems:

A detailed power analysis of a programmable DSP processor and an integrated RISC and DSP processor was described in the papers by Bajwa and Kojima et al [2,3]. The authors have shown a module-wise breakdown of power used in the different blocks. Contrary to many opinions it was found that the bus power is significantly smaller compared to the data path. It was also shown in this paper how a simple switch of the multiplier inputs (applicable to Booth encoded multipliers only) can reduce multiplier power by 4-8 times. Instruction fetch and decode contribute a significant portion of the power in these designs and since signal processing applications tend to spend a very large portion of their dynamic execution time executing loops, simple buffering schemes (buffers/caches) help reduce power by up to 25% [10].

Table 3. Comparison of Performance/Power and 1/Energy*Delay for representative RISC microprocessors[8,9]

Feature	Digital 21164	MIPS 10000	PwrPC 620	HP 8000	Sun Ultra-Sparc
Freq	500 MHz	200 MHz	200 MHz	180 MHz	250 MHz
Pipeline Stages	7	5-7	5	7-9	6-9
Issue Rate	4	4	4	4	4
Out-of-Order Exec.	6 lds	32	16	56	none
Register Renam. (int/FP)	none/8	32/32	8/8	56	none
Transistors/ Logic transistors	9.3M/ 1.8M	5.9M/ 2.3M	6.9M/ 2.2M	3.9M*/ 3.9M	3.8M/ 2.0M
SPEC95 (Intg/FlPt)	12.6/18.3	8.9/17.2	9/9	10.8/18.3	8.5/15
Power	25W	30W	30W	40W	20W
SpecInt/ Watt	0.5	0.3	0.3	0.27	0.43
1/Energy*Delay	6.4	2.6	2.7	2.9	3.6

In Fig. 2, the power breakdown is shown for the integrated RISC+DSP processor. For the benchmarks considered, which are kernels for typical DSP applications, the CPU functions as an instruction fetch, instruction decode and address generation unit for the DSP. Hence the variability in its power is less.

Similarly, the power consumed in the memories is quite high (in spite of their being low power, segmented bit-line memories) and shows little variation. In the case of the DSP the power variation is more and is data dependent. The interconnect power (INTR) represents the top level interconnect power which includes the main busses (three data and three address) and the clock distribution network at the top level. Clock power alone contributes between 30 and 50% of the total power consumption depending on system load.

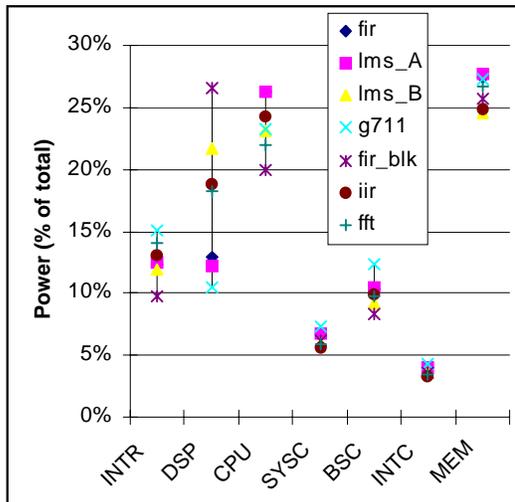
Table 4. A difference in power-performance factor resulting from increasing the size of cache[8,9]

Feature	401	403	Difference
Frequency	50MHz	66MHz (50MHz)	close
CMOS Process	0.5u 3-metal	0.5u 3-metal	same
Cache Total	2K-I / 1K-D	16K-I / 8K D	8x
FPU	No	No	same
MMU	No	Yes	
Bus Width	32	32	same
Transistors	0.3 Million	1.82 Million	600%
MIPS	52	81 (61)	+56% (+17%)
Power	140mW	400mW (303mW)	+186% (+116%)
MIPS/Watt	371	194	-91%

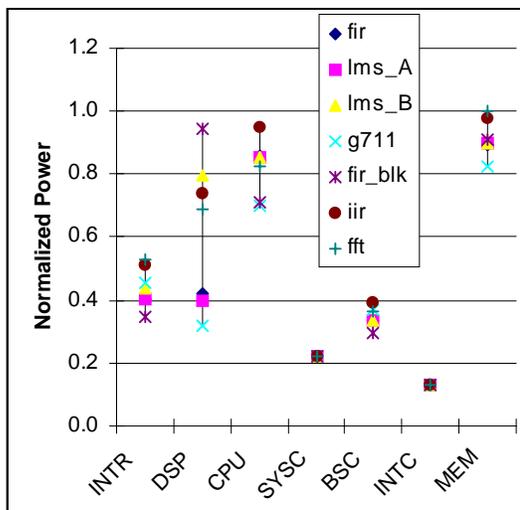
Table 5. The effect of increasing the cache size of PowerPC architecture[8,9]

Feature	604	620	Difference
Frequency	100MHz	133MHz (100MHz)	same
CMOS Process	0.5u 4-metal	0.5u 4-metal	same
Cache Total	16K+16K Cache	64K	~double
Load/Store Unit	Yes	Yes	same
Dual Intgr Unit	Yes	Yes	same
Reg- Renaming	Yes	Yes	same
Peak Issue	4 Instructions	4 Instructions	same
Transistors	3.6 Million	6.9 Million	+92%
SPECint92	160	225 (169)	+6%
SPECfp02	165	300 (225)	+36%
Power	13W	30W (22.5W)	+73%
Spec/Watt	12.3 / 12.7	7.5 / 10	-64%

The characteristics of embedded systems are quite different from those of desktop systems. For one, cost is a much more acute issue. Secondly, the computational load is a smaller well-defined set of tasks. For real-time signal processing applications, throughput behavior is typically more critical than minimum response time. These constraints dominate the design decisions. In many instances the cost of packaging is comparable to the cost of the die and using a more expensive package albeit with better heat dissipation capabilities is not an option. In the mobile arena, battery life and heat dissipation in compact designs (constricted space reduces airflow and hence the capacity to disperse heat) put downward pressure on power consumption of these processors. Depending on the application domain there are two broad approaches.



(a)



(b)

Fig. 2 Module-wise breakdown of the chip power consumption for the kernel benchmarks for the integrated RISC+DSP processor, (a) as a percentage of the total (b) normalized

Throughput and real-time constraints typically lead to more balanced systems, as in the case of DSPs (Harvard architecture, processor speed equal to bus speed). Balance here is a reference to a balance between Throughput and real-time constraints typically lead to more balanced systems, as in the case of DSPs (Harvard architecture, processor speed equal to bus speed). Balance here is a reference to a balance between processing speed, bandwidth and memory. Portable computing devices such as PDAs and handheld PCs form the other application domain, one in which the processors see a load similar to that of desktop systems in many respects. The StrongARM drops the processor core's clock frequency to be equal to that of its bus' clock frequency when it makes accesses off-chip thereby curtailing its power allowing its MIPS/Watt rating to scale.

Benchmarks are fraught with controversy and in the case of embedded systems where MIPS numbers are

based on Dhrystones, it is especially meaningless. The Dhrystone suite can fit in roughly 4KB of memory. This makes the disparity or lack thereof between processor speeds and bus speeds noteworthy.

The biggest impact on performance/power is process technology. The StrongARM, which is at the high end of embedded and low power processors, benefits from DEC's process technology (same as the one used for the Alpha chips) and a full custom design. This is atypical of embedded processor design. As recently as 1.5 years ago, the SA-110 was available in 0.35 micron technology and 2/3.3V (core/IO). All of its competitors were available in technologies ranging from 0.5 to 1 micron and voltages between 3.3V and 5V. This is changing but the SA-110 and the SA-1100 have been able to maintain their leading position as low power processors by aggressively reducing the core's voltage (1.35V for the SA-1100), circuit techniques and edge triggered flip-flops. A threshold voltage of 0.35 has allowed a much lower operating voltage. Most other embedded processors have had higher threshold voltages and hence, correspondingly higher operating voltages. Over the next year or two embedded processors with lower threshold voltages and dual threshold designs will become more standard.

The ARM9TDMI which has adopted a SA-110-like, five-stage pipeline, as opposed to ARM's traditional three-stage design, and a Harvard architecture illustrates the advantages of a more balanced design and can now be clocked at 150 MHz at sub-watt power levels. Better task partitioning is possible in embedded systems, due to the applications requiring a small set of predictable tasks to be performed, allowing unused hardware to be shutdown. In DSPs, control overheads are minimized and the data-path power and activity dominates. In desktop processors by contrast the control power almost drowns out the variations in the data-path power [4]. Power analysis of DSPs and simple RISC processors show two main sources of power the data-path units (multiply-accumulate units) and memory or cache (Fig.2.).

Conclusion

The conclusion from the studies presented is that the best power-performance is obtained if the architecture is kept simple thus allowing improvements to be achieved by technology. In the other words, the architecture should not stay in the way of technology and whenever this is not the case we will experience a decrease in power efficiency.

The second finding that goes contrary to the common knowledge is that we should seek improvements via simple design but increasing the clock frequency rather than keeping the frequency of operation low and increasing the complexity of the design.

The current processors today have reached their limit in terms of power. Digital 21264 is an example of a processor which had a potential of higher operating frequency but had to lower it (to 600MHz) in order to keep the power contained. This situation was first reached by Digital "Alpha" processor but it is soon to be reached by all of the others.

More specialized systems, used in signal processing applications can benefit from re-configurable data-path designs. The main advantage is to reduce the clock and control overhead by mapping loops directly onto the re-configurable data-path.

Applications in signal processing where stream data- or block data-processing dominates it makes sense to configure the data-path to compute algorithm specific operations. The cost of configuration can be amortized over the data block or stream. Aggressive use of chaining (as in vector processing) can be used to reduce memory accesses resulting in designs that may be called re-configurable vector pipelines. Embedded architectures can, in the future, be expected to employ all or some of these techniques

Acknowledgment

Contribution of Dr. Raminder Bajwa of Hitachi Semiconductor Research Laboratories as well as support received from Hitachi Ltd. is greatly appreciated.

References

1. T. Kuroda and T. Sakurai "Overview of Low-Power ULSI Circuit Techniques", IEICE Trans. Electronics, E78-C, No.4, April 1995, pp.334-344, INVITED PAPER, Special Issue on Low-Voltage Low-Power Integrated Circuits.
2. R. S. Bajwa, N. Schumann and H. Kojima, "Power analysis of a 32-bit RISC integrated with a 16-bit DSP", SLPED'97.
3. H. Kojima, et al, "Power Analysis of a Programmable DSP for Architecture/Program Optimization", Proceedings of the 1995 Low-Power Symposium.
4. V. Tiwari, S. Malik and A. Wolfe, "Power Analysis of Embedded Software: A First Step Towards Power Minimization", IEEE Trans. on VLSI Systems, 2(4):437-445, Dec. 1994.
5. Y. Sasaki, et al, "Multi-Level Pass-Transistor Logic for Low-Power ULSIs", Proceedings of the 1995 Low-Power Symposium.
6. C. Tan, et al, "Minimization of Power in VLSI Circuits Using Transistor Sizing, Input Ordering, and Statistical Power Estimation", Proceedings of the International Workshop on Low-Power Design, 1994.
7. M. Horowitz, et al, "Low-Power Digital Design", Proceedings of the 1994 IEEE Symposium on Low-Power Electronics, 1994.
8. L. Gwennap, "RISC on the Desktop: A Comprehensive Analysis of RISC Microprocessors for PCs, Workstations, and Servers", MicroDesign Resources, 1994.
9. Microprocessor Report, several issues, MicroDesign Resources, Sebastopol, California.
10. R. S. Bajwa et al, "Instruction Buffering to Reduce Power in Processors for Signal Processing", IEEE Trans. VLSI Systems, 5(4):417-424, December 1997.