Early-Stage Definition of LPX: a Low Power Issue-Execute Processor

P. Bose, D. Brooks, A. Buyuktosunoglu,
P. Cook, K. Das,² P. Emma, M. Gschwind,
H. Jacobson, T. Karkhanis,³ P. Kudva, S. Schuster,
J. E. Smith⁴, V. Srinivasan, V. Zyuban

IBM T. J. Watson Research Center Yorktown Heights, NY

Abstract

We present the high-level microarchitecture of LPX: a low-power issue-execute processor prototype that is being designed by a joint industry-academia research team. LPX implements a very small subset of a RISC architecture, with a primary focus on a vector (SIMD) multimedia extension. The objective of this project is to validate some key new ideas in power-aware microarchitecture techniques, supported by recent advances in circuit design and clocking.

1. Introduction

Power dissipation limits constitute one of the primary design constraints in future high performance processors. Also, depending on the thermal time constants implied by the chosen packaging/cooling technology, on-chip *power-density* is a more critical constraint than overall power in many cases. In current CMOS technologies, dynamic ("switching") power still dominates; but, increasingly, the static ("leakage") component is threatening to become a major component in future technologies [1]. In this paper, we focus primarily on the dynamic component of power dissipation.

Current generation high-end processors like the IBM POWER4TM [2,3] are performance-driven designs. In POWER4, power dissipation is still comfortably below the 0.5 watts/sq. mm. power density limit afforded by the package/cooling solution of choice in target server markets. However, in designing and implementing future processors (or even straight "remaps") the power (and especially the power-density) limits could become a potential "show-stopper" as transistors shrink and the frequency keeps increasing.

Techniques like clock-gating (e.g. [4, 5]) and dynamic size adaptation of on-chip resources like caches and queues (e.g. [6-10, 25, 27]) have been either used or proposed as methods for power management in future

D. Albonesi, S. Dwarkadas Univ. of Rochester, NY

processor cores. Many of these techniques, however, have to be used with caution in server-class processors. Aspects like reliability and inductive noise on the power supply rails (Ldi/dt) need to be quantitatively evaluated prior to committing a particular gating or adaptation technique to a real design.

Another issue in the design of next generation, power-aware processors, is the development of accurate power-performance simulators for use in early-stage design. University research simulators like Wattch [11] and industrial research simulators like Tempest [12] and PowerTimer [13] have been described in the recent past; however their use in real design environments is needed to validate the accuracy of the energy models in the context of power-performance tradeoff decisions made in early design.

In the light of the above issues, we decided to design and implement a simple RISC "sub-processor" test chip to validate some of the key new ideas in adaptive and gated architectures. This chip is called: LPX, which stands for low-power issue-execute processor. This is a research project, with a goal of influencing real development groups. LPX is a joint university-industry collaboration project. The design and modeling team is composed of 10-12 part-time researchers spanning the two groups (IBM and University of Rochester) aided by several graduate student interns and visiting scientists recruited from multiple universities to work (part-time) at IBM. LPX is targeted for fabrication in a CMOS 0.1 micron high-end technology. RTL (VHDL) simulation and verification is scheduled for completion in 2002. Intermediate circuit test chips are in plan (mid- to late 2002) for early validation of the circuit and clocking support. LPX chip tapeout is slated for early 2003. In this paper, we present the microarchitecture definition with preliminary simulation-based characterization of the LPX prototype.

We summarize the goals of the LPX project as follows:

To understand and assess the true worth of a few key ideas in power-aware microarchitecture design through simulation and eventually via direct hardware measurement. Based on known needs in real products of the future, we have set a target of average *power density* reduction by at least a factor of 5, with no more than 5 % reduction in

^{1, 2, 3:} A. Buyuktosunoglu, K. Das and T. Karkhanis were summer interns at IBM; they are currently back at Univs. of Rochester, Michigan and Wisconsin respectively.

^{4:} Prof. J. E. Smith, Univ. of Wisconsin, was on a sabbatical leave at IBM during the initiation of this project.

architectural performance (i.e. instructions per cycle or IPC).

- To quantify the instantaneous power (current) swings incurred by the use of the adaptive resizing, throttling and clock-gating ideas that are used to achieve the targeted power reduction factors in each unit of the processor.
- To use the hardware-based average and instantaneous power measurements for calibration and validation of energy models used in early-stage, power-performance simulators.

Clearly, what we learn through the "simulation and prototyping in the small" experiments in LPX, will be useful in influencing full-function, power-efficient designs of the future. The calibrated energy models will help us conduct design space exploration studies for high-end machines with greater accuracy. In this paper, we limit our focus to the microarchitectural definition process, with related simulation-based result snapshots, of the LPX prototype.¹

2. Background: Power-Performance Data

In an out-of-order, speculative super scalar design like each of the two cores in POWER4, a large percentage of the core power in the *non-cache* execution engine is spent in the instruction window or issue queue unit [3, 8-10]. Figure 1(a) shows the relative distribution of power across the major units within a single POWER4 core. Figure 1(b) zooms in on the instruction sequencing unit that contains the various out-of-order issue queues and rename buffers.

Figure 2 shows the power *density* across some of the major units of a single POWER4 core. The power figures are non-validated *pre-silicon* projections based on unconstrained (i.e. without any clock-gating assumptions) "average/max" power projections using a circuit-level simulation facility called CPAM [22]. (Actual unit-wise



Fig. 1(a). Relative (unconstrained) unit-wise power



Fig. 1(b). Power breakdown: ISU

power distribution, with available conditional clocking modes enabled, are not shown). This tool allowed us to build up unit-level power characteristics from very detailed, macro-level data. Here, the activity (utilization) factors of all units are assumed to be 100 % (i.e. worst case with no clock-gating anywhere); but average, expected input data switching factors (based on representative test cases run at the RTL level, and other



Fig. 2. Unconstrained power density profile

heuristics) are assumed for each circuit macro. Such average macro-level input data switching factors typically range from 4-15 %. (From Figure 2, we note that although on a unit basis, the power density numbers are under 0.5 watts/sq. mm., there are smaller hotspots, like the integer FX issue queue within the ISU that are above the limit in an unconstrained mode).

(Legend for Figs. 1-2: IDU: instruction decode unit; FXU: fixed point unit; IFU: instruction fetch unit; BHT: branch history table; ISU: instruction sequencing unit; LSU: load-store unit: includes L1 data cache; FPU: float-ing point unit).

Another class of data that we used was the performance and utilization information obtained from pre-silicon performance simulators. Figure 3 shows the relative "active/idle" barchart plot across some of the

¹ Note that LPX is a research test chip. It is not intended to be a full-function, production-quality microprocessor. At this time, LPX is not directly linked to any real development project.

major units for a POWER4-like pre-silicon simulation model. The data plotted is for a commercial TPC-C trace segment. This figure shows, for example, that the instruction fetch unit (IFU) is idle for approximately 47 % of the cycles. Similar data, related to activities within other units, like issue queues and execution unit pipes were collected and analyzed.



Figure 3. Unit-wise utilization stack (TPC-C trace)

3. Areas of Focus in Defining the LPX processor

Based on microarchitecture level and circuit simulation level utilization, power and power-density projections, as above, we made a decision to focus on the following aspects of a super scalar processing core in our LPX test chip:

Power-Efficient, Just-in-Time Instruction Fetch:

Here, we wanted to study the relative advantages of conditional gating of the ifetch function, with a goal of saving power without appreciable loss of performance. The motivation for this study was clearly established after reviewing data like that depicted in Figures 1 and 2. In simulation mode, we studied the benefit of various hardware heuristics for determining the "gating condition" [14, 23, 15] before fixing on a particular set of choices (being reported in detail in [16]) to implement in LPX. Our emphasis here is on studying ifetch gating heuristics that are easy to implement and test, with negligible added power for the control mechanism.

Adaptive Issue Queue:

The out-of-order issue queue structure inherent in today's high-end super scalar processors is a known "hot-spot" in terms of power dissipation. The data shown in Figures 1, 2, and also corroborative data from other processors (e.g. [5]), makes this an obvious area to focus on. In LPX, our goal is also to compare the achieved power savings with a fixed issue queue design, but with fine-grain clock-gating support, where the valid-bit for each issue queue entry is used as a clock-gating control. A basic issue in this context is the extra power that is spent due to the presence of out-of-order execution modes. Is the extra power spent worth the performance gain that is achievable? We wish to understand the fundamental power-performance tradeoffs in the design of issue queues for the next generation processors. Again, simplicity of the adaptive control and monitoring logic is crucial, especially in the context of the LPX prototype test vehicle.

Locally Clocked Execution Pipeline:

Based on the data shown in Figures 1 and 2, a typical, multi-stage complex arithmetic pipeline is also a high power-density region within the chip. We wish to study the comparative benefit of alternate conditional clocking methods proposed in ongoing work in advanced circuit design groups ([4, 17, 18]). In particular, we wish to understand: (a) the benefit of simple valid-bit-based clock-gating in a synchronously clocked execution unit; and (b) the added power-savings benefit of using a locally asynchronous arithmetic unit pipeline, within a globally synchronous chip. The asynchronously clocked pipeline structure is based on the IPCMOS circuit technology previously tested in isolation [17] by some in our research team. Such locally clocked methods offer the promise of low power at high performance, with manageable inductive noise (Ldi/dt) characteristics. In LPX, we wish to measure and validate these expectations as the IPCMOS pipe is driven by data in real computational loop kernels.

Power-Efficient Stalling of Synchronous Pipelines:

In the synchronous regions of the design, we wish to quantify the amount of power that is consumed by pipeline stall (or "hold/recirculation") conditions. Anticipating (from circuit simulation coupled with microarchitectural simulation data) such wastage to be significant, we wish to experiment with alternate methods to reduce or eliminate the "stall energy" by using a novel circuit technique called interlocked synchronous pipelines (ISP) that was recently invented by some members of our team [18].

Thus, a basic fetch-issue-execute super scalar processing element (see sections 4 and 5) was decided upon as the study vehicle for implementation by our small research team. The goal is to study the power-performance characteristics of *dynamic adaptation*: in microarchitectural terms as well as in clocking terms with the target of achieving significant power (and especially, power *density*) reduction, with acceptable margins of IPC loss.

4. Tuning the Microarchitecture

In this section, we outline the methodology adopted for defining the range of hardware design choices to be studied in the LPX testchip. Since we are constrained by the small size of our design team, and yet the ideas explored are targeted to influence real, fullfunction processor designs, we adopted the following general method.

Figure 4 shows the iterative method used to decide what coarse-level features to add into the LPX test chip, starting from an initial, baseline "bare-bones" fetch-issue-execute model.



Fig. 4. Methodology for fine-tuning the LPX microarchitecture

- A given, power-efficient microarchitectural design idea is first simulated in the context of a realistic, current generation super scalar processor model (e.g. POWER4-like microarchitectural parameters) and full workloads (like SPEC and TPC-C) to infer the power-performance benefit. Once a basic hardware heuristic is found to yield tangible benefit *in other words, a significant power reduction, at small IPC impact* - it is selected for possible implementation in LPX.
- A detailed, trace-driven, cycle-by-cycle simulator for the baseline LPX processor is coded to run a set of application-based and synthetic loop test cases designed to test and quantify the LPX-specific power-performance characteristics of the candidate hardware power-saving feature. In order to get a measurable benefit, it may be necessary to further simplify the heuristic, or augment the microarchitecture minimally to create a new baseline. Once the power-performance benefit is deemed significant, we proceed to the next candidate idea.

In this paper we mainly focus on (b) above: i.e. understanding the fundamental power-performance tradeoff characteristics, using a simple, illustrative loop test case. However, we also refer briefly to example, fullmodel super scalar simulation results to motivate the choice of a particular hardware heuristic.

Energy models used

The LPX cycle-by-cycle simulator used to microarchitectural early stage poweranalyze performance tradeoffs has integrated energy models, as in the PowerTimer tool [13]. These energy models were derived largely from detailed, macro-level energy data for POWER4, scaled for size and technology to fit the requirements of LPX. The CPAM tool [22] was used to get this data for most of the structures modeled. Additional experiments were performed at the circuit simulation level, to derive power characteristics of newer latch designs (with and without clock- and stall-based gating). The energy model-enabled LPX simulator is systematically validated using specially architected testcases. Analytical bounds modeling is used to generate bounds on IPC and unit-wise utilization (post-processed to form power bounds). These serve as reference "signatures" for validating the power-performance simulator. Since the LPX design and model are still evolving, validation exercises must necessarily continue throughout the high-level design process. Details of the energy model derivation and validation are omitted for brevity.

5. High-Level Microarchitecture of LPX

Figure 5 (placed at the very end of this paper) shows a very high-level block diagram of the baseline LPX processor that we started with before further refinement of the microarchitectural features and parameters through a simulation-based study. The function and storage units shown in dashed edges are ones that are modeled (to the extent required) in the simulation infrastructure, but are not targeted for implementation in the initial LPX design. The primary goal of this design is to experiment with the fetch-issue-execute pipe which processes a basic set of vector integer arithmetic instructions. These instructions are patterned after a standard, 4x32 bit SIMD multimedia extension architecture [21] but, simplified in syntax and semantics.

The "fetch-and-issue" sub-units act together as a *producer* of instructions, which are *consumed* by the "execute" sub-unit. The design attempts to balance the dynamic complexity of the producer-consumer pair with the goal of maximizing performance, while minimizing power consumption.

The basic instruction processing pipeline is illustrated in Figure 6. The decode/dispatch/rename stage, which is shown as a lumped, dummy dispatch unit in Figure 5, is actually modeled in our simulator as an m-stage pipe, where m=1 in the nominal design point. The nominal VFXU execute pipe is n=4 stages deep. The LSFX execute pipe is p=2 stages (in infinite cache mode) and p=12 stages when a data cache miss stall is injected using the stall control registers (Figure 5); in particular, using a miss-control register (MCR).

One of the functional units is the scalar FXU (a combined load-store unit and integer unit, LSFX) and the other is the vector integer unit (VFXU). The VFXU execution pipe is multi-cycle (nominally 4 cycles). The LSFX unit has a 1-cycle pipe plus (nominally) a 1-cycle (infinite) data cache access cycle for loads and stores. At the end of the final execution stage, the results are latched on to the result bus while the target register tags are broadcast to the instructions pending in the issue queue.



Fig. 6. LPX (simulator) pipeline stages

As a substitute for instruction caching, LPX uses a loop buffer in which a loop (of up to 128 instructions) is pre-loaded prior to processor operation. The loaded program consists of pre decoded instructions, with inline explicit specifiers of pre renamed register operands - in full out-of-order mode of execution. This avoids the task of designing explicit logic for the instruction decode and rename processes. LPX also supports an "in-order" mode, without register renaming as the lowest performance design point for our tradeoff experiments.

The instructions implemented in LPX are listed below in Table 1. For the most part, these are a set of basic vector (SIMD) mode load, store and arithmetic instructions, following the general semantics of a PowerPCTM VMX (vector multimedia extension) architecture [21]. There are a few added scalar RISC (PowerPC-like) instructions to facilitate loading and manipulation of scalar integer registers required in vector load-store instructions. The (vector) load and store instructions have an implied "update" mode in LPX, where the scalar address base register is auto-incremented to hold the address of the next sequential array data in memory.

Table 1. LPX Instruction Set				
	Example Syntax	Semantics		
Vector	VLD vr1, r2,	Load vr1. Scalar		
Load	0x08	base address regis-		
		ter: r2		
Vector	VST vr1, r2, 0x08	Store vr1.		
Store				
Vector	VADD vr1, vr2,	vr1 < vr2 + vr3		
Add	vr3			
Vector Sub, Mul, Div instructions: similar to VADD				
above				
Scalar	LD r1, r2, 0x08	Load scalar reg r1		
Load				
Scalar Inc	INC r1	Increment r1		
		(scalar)		
Scalar Dec	DEC r1	Decrement r1		
Cond.	BC +-0x08	Branch		
Branch		conditional		
		(PC relative jump)		
Uncond.	BR +-0x08	Branch		
Branch		unconditional		

1 1 DV 1

C

6. Examples: LPX Microarchitecture Analysis

In this section, we illustrate the use of simple loop-based test cases in understanding the basic powerperformance trade-offs of adaptive structures and clocking mechanisms that were chosen for study in LPX. The challenge is to determine the nominal sizes, adaptation windows and (in each case) a simple "monitor-andcontrol" mechanism that is appropriate in the context of building a small prototype engine, like LPX.

We started with the simplest baseline, where ideal cache effects were modeled, by architecting a single-stage LSFX pipe unit; but, later we had to augment the specification to include a variable-length LSFX pipe, to simulate data cache miss latency. In the absence of real cache hardware (correspondingly, real cache hit/miss code in the simulator), we architect for programmable "miss" scenarios via a user-loadable miss control register (MCR). Details of how this works in the real hardware are not discussed in this paper. For brevity, we only show a few example tradeoff analysis examples limited to the infinite (ideal) cache scenario.

As described before in section 4 (see Figure 4), each candidate power reduction idea is analyzed in the "large" (i.e. using a general out-of-order super scalar simulator) to ensure potential benefit. Then, a simpler hardware heuristic is used for trial and measurement "in the small" within the LPX simulation tool kit.

LPX experiments: an example loop test-case: **vect_add**

We use a simple "vector add" loop trace, formed by execution of the following loop, to illustrate LPX tradeoff experiments:

->	VLD	vr1,	r2 (0x4)
	VADD	vr4,	vr1, vr6
	VLD	vr6,	r2 (0x8)
	VADD	vr4,	vr4, vr6
	VST	vr4,	r3 (0x8)
	DEC	r7	
	BRZ	r7,	-0x7

The baseline LPX model parameters were fixed as follows, after initial experimentation. Instruction fetch (ifetch) bandwidth is up to four instructions/cycle, with no fetch beyond a branch on a given cycle. The instruction fetch buffer size is four instructions; dispatch bandwidth (into the issue queue) is up to two instructions/cycle; issue bandwidth (into the execution pipes) is up to two instructions/cycle; and, completion bandwidth is also two instructions/cycle. Fetch and dispatch is in-order and issue can be in-order or out-oforder (switchable); instructions finish out-of-order. (LPX does not model or implement in-order completion for precise interrupt support using reorder buffers).

Conditional ifetch

Figures 7(a,b) show a snapshot of analysis data from a typical 4-way, out-of-order super scalar processor model. The data reported is for two benchmarks from the SPECint2000 suite. It shows that the ifetch stage/buffer, the front-end pipe and the issue queue/window can be idle for significant fractions of the program run. These are cycles where power can be saved by valid-bit-based clock-gating. In addition, the fraction of cycles that are wasted by useful (but stalled) instructions and by incorrectly fetched speculative instructions can also be significant. Gating off the ifetch process using a hardware heuristic to compute the gating condition, is therefore a viable approach to saving energy.

For LPX, we wish to experiment with the simplest of such heuristics, that are easy to implement. The basic method used is to employ the "stall" or "impending stall" signals available from "downstream"



Fig. 7(a). Idle, speculative and stall waste: AMM



Fig. 7(b). Idle, speculative and stall waste: GCC

consumer units to throttle back the "upstream" producer (ifetch). Such stall signals are easy to generate and are usually available in the logic design anyway. Figures 8(a,b) show results from an illustrative use of conditional ifetch while simulating the vect_add loop trace.

We use the following simple hardware heuristic for determining the ifetch gating scenario. When a "stall" signal is asserted by the instruction buffer (e.g. when the ibuffer is full) the ifetch process is naturally inhibited in most designs; so this is assumed in the baseline model. However, additional power savings can be achieved by retaining the "ifetch-hold" condition for a fetch-gate cycle window, GW, beyond the negation of the ibuffer stall signal. Since the ibuffer was full, it would take a while to drain it; hence ifetch could be gated off for GW cycles. Depending on the size of the ibuffer, IPC would be expected to drop off to unacceptable levels beyond a certain value of GW; but increasing GW is expected to reduce IFU (instruction fetch unit) and overall chip power.

Adaptive Issue Queue

Figure 9 shows a snapshot of our generalized simulationbased power-savings projection for various styles of out-of-order issue queue design. (An 8-issue, super scalar, POWER4-like research simulator was used). These studies showed potential power savings of more



CPI of baseline in-order (IO) = 3.00CPI of all the other out-of-order (OO) = 2.29Fig. 8(a). Conditional ifetch in LPX: power reduction



Fig. 8(b). Cond. ifetch in LPX: CPI vs. gating wind

than 80 % in the issue queue with at most 2-3 % hit in IPC on the average. However, the best power reductions were for adaptive and banked CAM/RAM based designs that are not easy to design and verify. For LPX, we started with a baseline design of the POWER4 integer issue queue [3], which is a latch-based design. It is structured as a 2-chunk structure, where in adaptive mode, one of the chunks can be shut-off completely (to eliminate dynamic *and* static power).

Figure 10 illustrates the benefit of using a simple, LPX-specific adaptive issue queue heuristic that is targeted to reduce power, without loss of performance; i.e. the size is adapted downwards only when "safe" to do so from a performance viewpoint; and the size is increased in anticipation of increased demand. (In the example data shown in this paper, we consider only the reduction of dynamic power via such adaptation). The adaptive issue queue control heuristic illustrated is simpler than proposed in the detailed studies reported earlier [8], for ease of implementation in the LPX context. The control heuristic in LPX is as follows:

if (current-cycle-window-issuecount < 0.5
last_cycle_window_issuecount) then
increase_size (* if possible*);
else decrease_size (* if possible *);</pre>

Discussion of results:

From Figure 8(a) we note that adding out-oforder (oo) mode to the baseline in-order (io) machine causes an IPC increase (CPI decrease) of 23.6 %, but with a 12.5 % overall power increase. The ISU, which contains the issue queue, increases in power by 27.5 %. So, from an overall power-performance efficiency viewpoint, the out-of-order (oo) mode does seem to pay off in LPX for this loop trace, in infinite cache mode. However, from a power-density "hot-spot" viewpoint, even this basic enhancement may need to be carefully



evaluated with a representative workload suite. Adding the valid-bit-based clock-gating (VB-CG) mode in the instruction buffer, issue queue and the execution unit pipes, causes a sharp decrease in power (42.4 % from the baseline oo design point). Adding a conditional ifetch mode, (with a cycle window W of 10 cycles over which ifetch is blocked after the ibuffer stall signal goes away) yields an additional 18.8 % power reduction, without loss of IPC performance. As the gating cycle window W is increased, we see a further sharp decrease in net power beyond W=10, but with IPC degradation. For the adaptive issue queue experiment (Fig. 10) shown, we see that a 8 % reduction in net LPX power is possible; but beyond an adaptation cycle window, AW of 1, a 11 % increase in CPI (cycles-per-instruction) is incurred. Thus, use of fine-grain, valid-bit based clock-gating is simpler and more effective than adaptive methods. Detailed results, combining VB-CG and adaptation will be reported in follow-up research.



Fig. 10. Adaptive issue queue experiment in LPX

Stall-Based Clock-Gating

As previously alluded to, in addition to valid-bitbased clock-gating in synchronous (an locally asynchronous) pipelines, LPX uses a mode in which an instruction stalling in a buffer or queue for multiple cycles is clock-gated off, instead of a recirculation-based, hold strategy often used in high performance processors. The stall-related energy waste is a significant fraction of queue/buffer power that can be saved if the stall signal is avaliable in time to do the gating. Carefully designed control circuits [18] have enabled us to exploit this feature in LPX. In this version of the paper, we could not include the experimental results that show the additional benefits of such stall-based gating. However, suffice it to say, with the addition of stall-based clock-gating, simulations predict that we are well within the target of achieving a factor of 5 reduction in power and power density, without appreciable loss of IPC performance. The use of a locally asynchronous IPCMOS execution pipe [17] is expected to increase power reduction even further. Detailed LPX-specific simulation results for these circuitcentric features, will be available in subsequent reports.

7. Conclusions and Future Work

We presented the early-stage definition of LPX: a low-power issue-execute processor prototype that is designed to serve as a measurement and evaluation vehicle for a few new ideas in adaptive microarchitecture and conditional clocking. We described the methodology that was used to architect and tune simple hardware heuristics in the prototype test chip, with the goal of drawing meaningful conclusions of use in future products. We presented a couple of simple examples to illustrate the process of definition and to report the expected power-performance benefits of the illustrated adaptive features.

The basic idea of *fetch-throttling* to conserve power is not new. In addition to work that we have already alluded to [14, 23, 15], Sanchez et al. [24] describe a fetch stage throttling mechanism for the G3 and G4 PowerPC processors. The throttling mode in the prior PowerPC processors was architected to respond to thermal emergencies. The work reported in [14, 23, 15] and the new gating heuristics described in this paper and in [16] are aimed at reducing average power during normal operation. Similarly, the adaptive issue queue control heuristics being developed for LPX are intended to be simpler adaptations of our prior general work [8].

We believe that the constraint of designing a simple test chip with a small design team forces us to experiment with heuristics that are easy to implement with low overhead. If some of these heuristics help create relatively simple power management solutions for a fullfunction, production-quality processor, then the investment in LPX development will be easily justified.

In addition to the adaptive microarchitecture principles alluded to above, the team is considering the inclusion of other ideas in the simulation toolkit; some of these remain candidates for inclusion in the actual LPX definition: at least for LPX-II, a follow-on design. The following is a partial list of these other ideas:

- Adaptive, power-efficient cache and register file designs: these were not considered for implementation in the initial LPX prototype, due to lack of seasoned SRAM designers in our research team. In particular, as a candidate data cache design for LPX-II, we are exploring ideas that combine prior energy-efficient solutions [6, 7, 25, 26] with recently proposed, high performance split-cache architectures ([19, 20]).
- Exploiting the data sparseness of vector/SIMD-mode execution, through hardware features that minimize clocking waste in processing vector data that contains lots of zeroes.
- Newer features that reduce static (leakage) power waste.
- Adding monitoring hardware to measure current swings in clock-gated and adaptive structures.

ACKNOWLEDGEMENTS

The authors are grateful to Dan Prener, Jaime Moreno, Steve Kosonocky, Manish Gupta and Lew Terman (all at IBM Watson) for many helpful discussions before the inception of the LPX design project. Special thanks are due to Scott Neely (IBM Watson), Michael Wang and Gricell Co (IBM Austin) for access to CPAM and the detailed energy data used in our power analysis. The support and encouragement received from senior management at IBM - in particular Mike Rosenfield and Eric Kronstadt - are gratefully acknowledged. The authors would like to thank Joel Tendler (IBM Austin) for his comments and suggestions during the preparation and clearance of this paper. Also, the comments provided by the anonymous reviewers were useful in preparing an improved final version; this help is gratefully acknowledged.

The work done in this project at University of Rochester is supported in part by NSF grants CCR-9701915, CCR--9702466, CCR--9705594, CCR-9811929, EIA--9972881, CCR--9988361 and EIA--0080124; by DARPA/ITO under AFRL contract F29601-00-K-0182; and by an IBM Faculty Partnership Award. Additional support, at University of Wisconsin, for continuation of research on power-efficient microarchitectures, is provided by NSF grant CCR-9900610.

REFERENCES

- 1. S. Borkar, "Design challenges of technology scaling," IEEE Micro, vol. 19, no. 4, July/Aug. 1999, pp. 23-29.
- C. Anderson et al., "Phyiscal design of a fourth-generation POWER GHz microprocessor," *ISSCC 2001 Digest of Tech. Papers*, Feb. 2001, p. 232.

- J. M. Tendler, S. Dodson, S. Fields, H. Le and B. Sinharoy, "POWER4 System Microarchitecture," <u>http://www-1.ibm.com/servers/eserver/pseries/hardware/w</u> <u>hitepapers/power4.pdf</u>, Oct. 2001.
- J. M. Rabaey and M. Pedram, ed., <u>Low Power Design</u> <u>Methodologies</u>, pp. 279-281, Kluwer, 1996.
- M. Gowan et al., "Power considerations in the design of the Alpha 21264 microprocessor," *Proc. ACM/IEEE Design Automation Conference (1998)*, pp. 726-731.
- D. H. Albonesi, "The inherent energy efficiency of complexity-effective processors," *Proc. ISCA Workshop* on *Power-Driven Microarchitecture*, June 1998.
- R. Balasubramonian, D. H. Albonesi, A. Buyuktosunoglu and S. Dwarkadas, "Memory hierarchy reconfiguration for energy and performance in general purpose architectures," *Proc.* 33rd. Int'l. Symp. on Microarchitecture, pp. 245-257, December 2000.
- 8. A. Buyuktosunoglu et al., "An adaptive issue queue for reduced power at high performance," *Proc. Workshop on Power-Aware Computer Systems, (PACS'00),* held in conjunction with ASPLOS, November 2000.
- D. Ponomarev, G. Kucuk and K. Ghose, "Dynamic allocation of datapath resources for low power," *Proc. Workshop* on Complexity-Effective Design (WCED-01), held in conjunction with ISCA, June/July 2001;
- D. Folegnani and A. Gonzalez, "Energy-effective issue logic," *Proc. ISCA-01*, pp. 230-239, June 2001.
- D. Brooks, V. Tiwari and M. Martonosi, "Wattch: a framework for architectural-level power analysis and optimization," *Proc. of the 27th Int'l. Symp. on Computer Arch.*, pp. 83-94, Vancouver, Canada, June 2000.
- A. Dhodapkar, C. H. Lim and G. Cai, "TEM²P²EST: a thermally enabled multi-model power/performance estimator," Proc. Workshop on Power-Aware Computer Systems (PACS'00), Cambridge, MA, Nov. 2000.
- D. Brooks, J.-D. Wellman, P. Bose and M. Martonosi, "Power-performance modeling and tradeoff analysis for a high-end microprocessor," *Proc. PACS'00*, Cambridge, MA, Nov. 2000.
- S. Manne, A. Klauser and D. Grunwald, "Pipeline Gating: Speculation Control for Energy Reduction," *Proc. Int'l. Symp. on Computer Architecture (ISCA)*, June 1998.
- 15. A. Baniasadi and A. Moshovos, "Instruction flow-based front-end throttling for power-aware high performance processors," *Proc. Int'l. Symp. on Low Power Electronics and Design* (ISPLED-01), August 2001.
- 16. T. Karkhanis et al. "Saving energy with just-in-time instruction delivery," under submission for publication.
- S. Schuster, W. Reohr, P. Cook, D. Heidel, M. Immediato and K. Jenkins, "Asynchronous interlocked pipelined CMOS operating at 3.3-4.5 GHz.," *ISSCC 2000 Digest of Tech. Papers*, pp. 292-293.
- H. Jacobson et al., "Synchronous interlocked pipelines," *Proc. ASYNC-2002* (to appear); prior version available as IBM Research Report RC 22239, October 2001.
- V. Srinivasan, Hardware Solutions to Reduce Effective Memory Access Time, Ph. D dissertation, Univ. of Michigan, Ann Arbor, Feb. 2001
- 20. V. Srinivasan, et al. "Recovering single cycle access of primary caches," submitted for publication.

- K. Diefendorff, P. K. Dubey, R. Hochsprung and H. Scales, "AltiVec[™] extension to PowerPC accelerates media processing," *IEEE Micro*, March/April 2000, pp. 85-95.
- J. S. Neely, H. H. Chen, S. G. Walker, J. Venuto and T. J. Bucelot, "CPAM: A Common Power Analysis Methodology for High Performance Design, *Proc. 9th Topical Meeting on Electrical Performance of Electronic Packaging*, Oct. 23-25, 2000, Scottsdale, AZ.
- D. Grunwald et al., "Confidence estimation for speculation control," *Proc. Int'l. Symp. on Comp. Arch., (ISCA), June* 1998.
- 24. H. Sanchez et al., "Thermal management systems for high performance PowerPC microprocessors," *Proc. COMPCON-97*, Feb. 1997.
- 25. D. H. Albonesi, "Selective cache ways: on-demand cache resource allocation," *Proc. 32nd. Int'l. Symp. on Microarchitecture (MICRO-32)*, pp. 248-259, Nov. 1999.
- K. Inoue et al. "Way-predicting set-associative cache for high performance and low energy consumption," *Proc. Int'l. Symp. on Low Power Electronics and Design* (*ISLPED*), pp. 273-275, Aug. 1999.
- S-H Yang et al., "An energy-efficient high performance deep submicron instruction cache," *IEEE Trans. on VLSI*, Special Issue on Low Power Electronics and Design, 2001.



Figure 5. LPX Processor: High-Level Block Diagram