

128 Bit Reconfigurable RISC Microprocessor Architecture

**A Phase 1 Proposal to Small Business Innovation
Research within the National Science Foundation (SBIR)
by**

Obsidian Technology

**in Association with Staff of the
Advanced Computer Architecture Group,
University of California at Irvine**

For questions about this project, please contact:

Robert Heaton

Obsidian Technology

37 Niguel Pointe Drive

Laguna Niguel

CA 92677

Phone: (714) 363-7982.

Fax: (714) 249-2167.

Table of Contents

D. Identification Significance of the Problem or Opportunity	5
E. Background and Technical Approach	6
Technology Background.....	6
Dynamically Programmable Instruction Unit.....	7
DPA Integration Issues	7
Instruction Set.....	8
DPA Code Example.....	9
Data Dependent ALU Timing Scheme.....	9
Thermally Constrained Clock Throttling.....	10
F. Phase I Research Objectives	11
G. Phase I Research Plan.....	12
Instruction Set Design	12
Architecture.....	12
DPA Porting Study	13
Develop Simple Assembler	13
Develop Test Code.....	13
System Simulation.....	13
Temperature Sensor.....	13
Document Results.....	13
H. Commercial Potential	13
Customer Needs	13
Customer Base	14
I. Principal Investigator and Senior Personnel.....	15
Robert Heaton	15
J. Consultants and Subawards.	15
Prof. Nader Bagherzadeh (University of California at Irvine)	15
Dr. Fadi J. Kurdahi (University of California at Irvine)	16
Prof. Nader Bagherzadeh. Consultant Letter.....	17
Dr. Fadi J. Kurdahi. Consultant Letter.....	18
NSF Budget Proposal Form 1030A.....	19
Appendix D. Supplemental: Budget Justification/Explanation Page.....	20
References.....	21
K. Equipment, Instrumentation, Computers, and Facilities.....	21
L. Current and Pending Support of the Principal Investigator and Senior Personnel	22
M. Equivalent or Overlapping Proposals to Other Federal Agencies.....	22
O. Prior SBIR Awards.....	22
Example APR128 Data Sheet.....	23

D. Identification Significance of the Problem or Opportunity

There is a significant difference between the requirements for general computing and the processing of very large data streams. While the general computing market remains dominated by the x86 architecture, the embedded processor field remains more open and architectures are more frequently selected on performance, cost, and power dissipation grounds rather than market predominance. Typical applications for embedded microprocessors include Laser printers, automotive control, IP packet routing, Personal Data Assistants (PDA), and games hardware. Hence the volume market for embedded processors is many times higher than that of the x86 volume.

Within the embedded processor field performance remains a key issue in applications which process large volumes of data or fast streams of data. This proposal seeks to produce a device which offers significant performance and cost benefits over today's 32 and 64 bit embedded CPUs and which will be scalable to arbitrary data bus widths without changes in the basic instruction set.

Whereas multi stage data caches have reduced the requirements for CPU to core memory bandwidth, it is




Figure 1. Data Flow Bottleneck in Typical Embedded RISC Application.

still a significant bottleneck, especially in streaming data or large data applications where the efficiency of a data cache is minimized. Such applications include IO processors, Local Area Network (LAN) switching and routing, and software based graphical processing (Eg printers and games). See Figure 1.

Another aspect of the opportunity related to the state of package development. Low cost 32 bit RISC processors were enabled by high volume

PLCC packages during the mid eighties. Today, significant improvement in device bonding and packaging technology permit low cost devices with over 200 pins to be developed.

Similarly the first generation of RISC devices employed 3μ double level metal CMOS technologies. Today's 0.35μ four layer metal CMOS technologies are approximately two orders of magnitude more dense. Hence although incremental additions have been made to first generation RISC architectures the scale of the underlying implementation technology advances suggest the value of re-architecting the single chip RISC to reach a new level of performance.

With the shrinking of process geometries the RISC CPU core has become a very small proportion of chip area. In 0.35μ CMOS MIPS and ARM architectures have 32 bit CPU core areas below 4mm^2 . This suggests that a 128 bit CPU architecture could be architected in core area of less than 16mm^2 , still a small area. Hence we feel that 128 bit CPU with a modest instruction cache, and a small DPA can be implemented and still meet the \$10 unit cost objective.

This technology background is reminiscent of the leap made between 8 bit CISC (Complex Instruction Set Computers) and 32 bit RISC processors during the eighties and we anticipate a similar value to the nation in terms of performance/cost ratio of consumer and industrial products. If successful, this research will support the commanding lead the US has in CPU manufacturing and further the computing goals of the NSF.

E. Background and Technical Approach

Technology Background

The current generation of embedded microprocessors are highly upgraded versions of RISC architecture developed more than a decade ago. Their small size, low cost, and low power has enabled them to be used in a wide variety of applications from Personal Digital Assistants (PDA) to high end single chip Network Computers.

This notwithstanding, it has been recognized by researchers for some time that the geometric progression in the density of silicon technologies will eventually offer the opportunity for some fundamental changes in embedded CPU architectures. Apart from evolutionary and Very Long Instruction Word (VLIW) approaches (for example [VIPER] ref 4) there are a number of ongoing research programs which attempt to provide closer association between memory and processing elements either by dispersing processing elements within conventional memory or by the use of RAM based Dynamically Reprogrammable Arrays. ([BRASS] ref 3, [MIT] ref 6). There are, however, significant practical barriers to the general use of such schemes due to the significant shift in the processing paradigm. In particular, it is difficult to leverage existing Operating System (OS) and applications software base.

Some researchers have focused on the integration of memory and processor within the same device. ([IRAM] Ref 5, Ref 10). We think this is a very promising approach for the following reasons:

1. One Giga Byte DRAM devices have recently been announced and demonstrated in the lab. These devices will not be available in volume for three or four years, but when that happens the memory access bottleneck will become so significant that some form of closer CPU-RAM coupling architecture will be mandated.
2. With an integrated CPU and large RAMs in one device large scalable arrays can be built allowing an “intelligent RAM” of arbitrary size to be constructed.

However to date only 32 bit processors and large RAM arrays have been integrated. We think very wide data path processors offer significant advantages for this approach since the data path width may be made to match the RAM column width as in Figure 2.

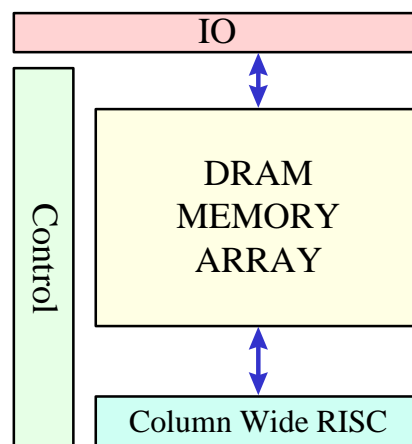


Figure 2. Integrated DRAM and CPU Chip.

Dynamically Programmable Instruction Unit

The technical approach taken in the proposal is a pragmatic melding of an existing conventional 32 bit RISC instruction set with a very wide data bus and the inclusion of a Dynamically Programmable Array (DPA) within the CPU data path. See Figure 3.

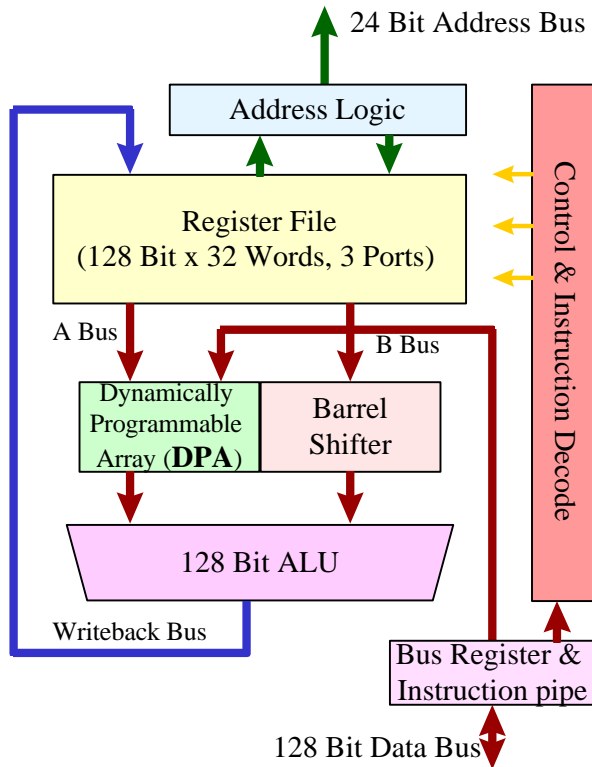


Figure 3. 128 Bit RISC CPU Architecture.

The DPA will be a custom designed multi plane FPGA VLSI structure which can effectively implement custom instructions in hardware. ([DPA] Ref 2, [MIT] ref 6) This permits application specific data processing instructions to be implemented directly in hardware.

During a custom instruction cycle primary read busses A and B are processed by the DPA at the same the B bus is being processed by the Shifter. Changing the DPA code footprint can be achieved within one clock cycle since a number of footprints are loaded into a multi plane RAM at boot time. The basic DPA element from the BAA 97-06 project is shown in Figure 4. As far as possible the cells and placement/routing software will be taken directly from the BAA 97-06 ARPA work or from [MIT] Ref 6.

Note that a very wide data bus has good potential for future segmentation in a similar fashion to the Intel MMX architecture.

DPA Integration Issues.

Although the DPA unit appears in the path of the A and B busses in much the same way that the Barrel Shifter appears in the B Bus (see Figure 3) its integration into the CPU is significantly more complex. The following issues must be resolved:

1. Code plane loading. The DPA is basically a RAM based FPGA with 2^n codes. Hence a method of loading these codes during operation is required. One way to achieve this is to memory map the code planes.
2. Dynamic timing. Since the DPA logic configuration changes with its code, DPA timing changes also. Hence it may change at each instruction boundary and may even be data dependent. The timing may be several CPU cycles, hence a DPA based instruction must be able to signal timing information to the controller.

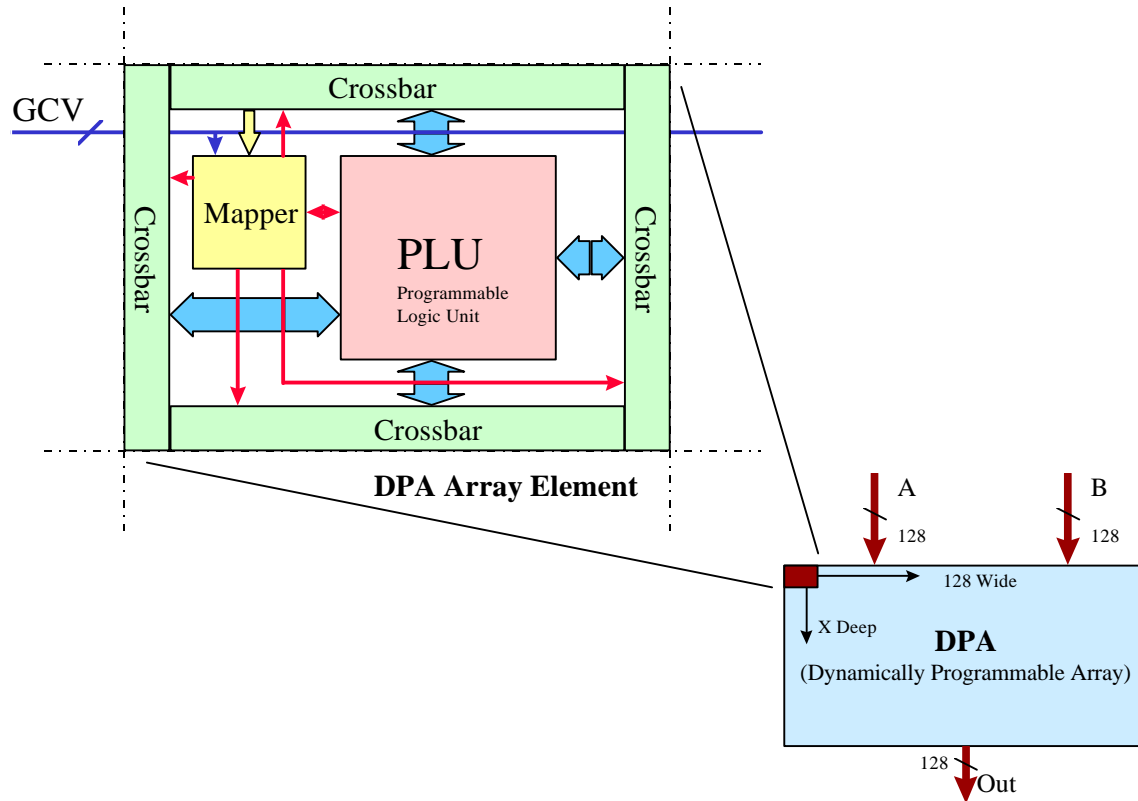


Figure 4. DPA Architecture.

Instruction Set

The introduction of a new processor without object code compatibility is a significant challenge. The porting of some operating systems can be done relatively easily, but the porting of a very large base of application code represents significant barriers of commercial viability and training as well as significant technical difficulties.

This proposal aims to mitigate these difficulties by maintaining object code compatibility with an existing, widely used, 32 bit embedded RISC while offering 128 bit data processing instructions which can be added selectively into embedded data intensive applications. Since the Principal Investigator, Robert Heaton, led the team which built the Acorn RISC Machine (ARM) during the early eighties, it is a natural choice for code compatibility. The ARM is a widely used embedded architecture licensed by a number of major Corporations such as Apple, Texas Instruments, Lucky Goldstar, Rockwell, NEC, Philips, Lucent, etc. It is used in a wide variety of applications from PDAs to networking products to data encryption and disk controllers. (For more information see <http://www.arm.com/>)

There are a number of challenges in combining a 32 bit instruction set with 128 bit data processing instructions which include:

1. Development of unique hooks into the 32 bit instruction space which do not inhibit 32 bit compatibility.
2. Maintaining compatibility while utilizing the full bandwidth of 128 bit instruction and data fetches.
3. Insuring that the 128 bit instruction mode can interface reasonably smoothly with 32 bit operating systems without significant loss of performance.

4. Appropriate handling 8 and 32 bit data operations within 128 bit data processing instructions.
5. Special instructions are required for loading and running the DPA.

DPA Code Example

The DPA will be controlled from the machine code in a similar way as a typical data processing instruction. Its operation may be combined with ALU and Shifter operation within one instruction. To take a simple example in a conventional CPU an addition of registers A and B with the result placed in register C may require an assembler phrase like: **ADD A, B → C;**

In the case of a DPA instruction a similar phrase can be used. Let us suppose our application is enhanced by having an instruction which reverses the order of bits in register A if the value in register B is 0xFFF3. In this case the system designer must:

1. Design the DPA code. This could be done with a schematically entered logic design or through logic synthesis.
2. Perform timing analysis for the code. In this case the logic is so simple that the operation is very likely to occur within a single cycle.
3. Simulate the operation of the new instruction in a system simulation tool.
4. Load the DPA code before its use in the instruction stream. Since the DPA code space is memory mapped this would normally just be a block memory move operation done once during initialization.
5. Call the code. We expect the simple case assembler syntax to be: “DPA <code #> <timing> <operand> <operand> → <destination>”

Hence for this example the assembler code could be: **DPA 1, 1, A, B → C;** I.e. take code plane one, with one cycle timing. Inputs are A and B, the result is placed in C.

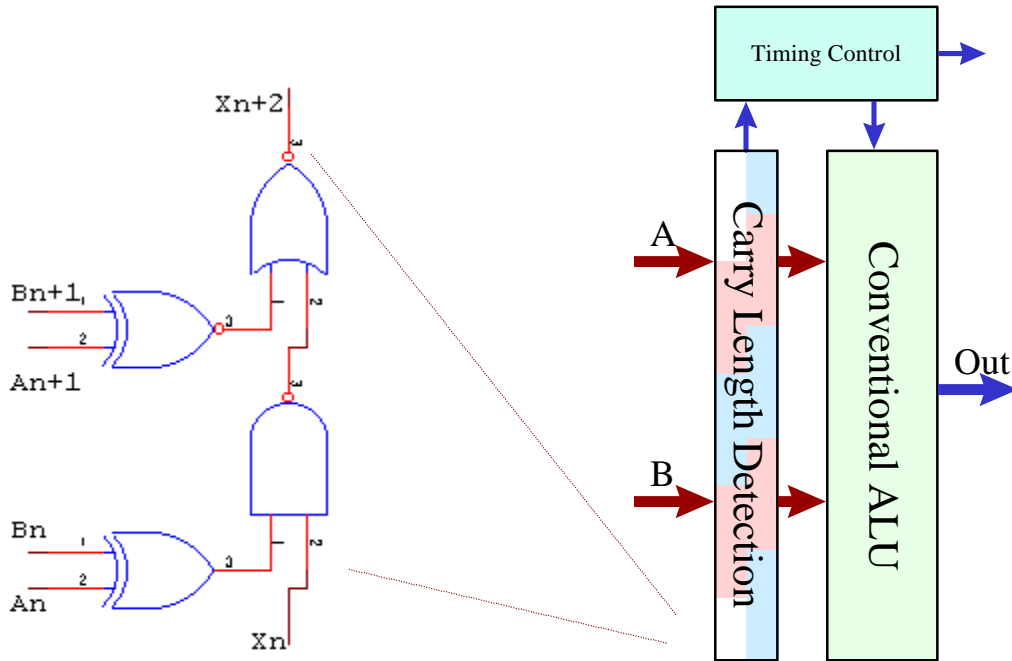
Data Dependent ALU Timing Scheme.

A difficulty with increased data path width (for example to 128 bits) has traditionally been that ALU delay times increase with path width due to the need to extend the carry chain length. Even with a complex nested carry look ahead scheme, 128 bit arithmetic will significantly slow down the worst case carry path. However, typical ALU delay times equivalent to, or less than, 32 bit circuits can be achieved by using a variable cycle time data dependent approach.

This approach makes use of low probability of a carry rippling for more than 32 bits. A carry length detection circuit extends the computation time when a carry will ripple for more than a specified number of bits occur within the 128 bit ALU. With random operands, we estimate time extension will occur about 1 in 2^{32} times.

In real programs of course the data is not random. A typical case that would cause time extension would be the relatively common situation of adding 1 to -1 in two's complement. However, this event occurs much less than 0.1% of the time and does not significantly degrade performance, especially if an optimizing compiler can avoid these operations.

For an ALU add operation bitwise operands A_n and B_n propagate carry timing only if $A_n \oplus B_n$ is true. This is because either a carry generate ($A_n \cdot B_n$) or a bitwise sum of zero ($A_n \cdot B_n = 0$) suppress the ripple



of

Figure 5. Carry Ripple Detection Primitive

the carry timing from bit N to bit $N+1$. Thus for each bit the carry ripple function $X_{n+1} = (A_n \oplus B_n) \cdot X_n$. Hence a simple circuit as shown in Figure 5 can be stacked to indicate a given ripple carry length. By overlapping a number of these stacks a circuit can estimate the time required for the addition to ripple through the ALU and adjust the CPU cycle timing accordingly. We will research the best overall statistical performance for a number of arrangements using this approach. Note that in Figure 5 the X_n to X_{n+1} delay is one gate delay and with appropriate asymmetrical gate transistor sizing this can be made to approach one gate delay. The previous work of Robert Heaton in the ARM1 fast ripple carry ALU ([VLSI] Ref 11) and the work of Prof. Nader Bagherzadeh ([RISC]Ref 1) is particularly relevant to this work.

Thermally Constrained Clock Throttling

Most embedded processors are able to go into an idle state when not actively processing in order to save power. Statistically, in many applications processors do useful processing in short bursts since there are typically delays for peripheral accesses or user input. Most CPU's, however, have to be designed into a thermal environment which assumes maximum environmental temperature, minimum power supply voltage, and 100% processor utilization. This generally requires increased cost packaging and/or specially cooling arrangements.

With the addition of a number of distributed critical delay measurement circuit (for example Figure 6) and an integrated temperature sensor when the die temperature reaches a preset limit the maximum clock rate is reduced. Since maximum ambient temperatures and continuous operation occur infrequently the device can be packaged in standard packaging without significantly reducing performance.

Clock throttling permits the chip to operate to its maximum thermal capability in any environment; whereas in conventional processors a significant safety margin is required to ensure functionality over a

range of operating environments. Hence, we propose an investigation into how this can be practically achieved.




Figure 6. Simple Method for Sensing Need for Throttling.

F. Phase I Research Objectives

Our long term goal is to create a viable, simple, low cost 128 bit embedded processor architecture for data processing applications for a manufacturing cost of below \$10. During Phase I we will demonstrate an architectural model, estimate relative performance, and relative cost. We will draw on a wealth of experience of the investigators and staff at UCI's Computer Laboratory. Specific objectives for Phase I are as follows:

1. Develop and document an instruction set for the 128 bit processor which achieves a good compromise between 32 bit backward compatibility and 128 bit data processing performance.
2. Develop and document an architecture for the processor, capturing it in a High level Description Language. (HDL)
3. Simulate the HDL model in a simple system environment with an extensive set of instruction cases.
4. Design CMOS junction temperature monitoring circuit and simulate with SPICE simulator.
5. A 128 bit dynamic ALU control circuit will be designed and simulated at a gate level.
6. Document the viability of importing the results of BAA 97-06 or [MIT] Ref 6 proposal into the basic architecture.

Upon conclusion of Phase I a firm foundation of technical viability will have been established which can be built on in Phase II during which a detailed prototype device and basic development tools will be developed.

Further research will be performed during Phase II to investigate device integrations which could profit from data path widths wider than 128 bits. In particular we are interested in high capacity RAM devices which have a very wide RISC processor data path connected directly to the RAM array columns. We think that array of such integrated [IRAM] devices can lead to the order of magnitude performance improvements sort by NSF's objectives. Hence offering an array-able intelligent memory device. It is our expectation that at the conclusion of Phase II Obsidian Technology will be able to obtain development funding from a base of potential licensers and/or venture capital sources.

G. Phase I Research Plan.

The member of the team documented in the budget will be assigned one or more of the tasks outlined below:

Instruction Set Design

This is the key activity in the proposal since it drives aspects of the architecture, DPA, assembler, and final documentation and will be performed in parallel with these activities. Instruction set will be defined by Nader Bagherzadeh, and a graduate student who will be responsible for consistency checking and documentation. Some of the key elements of the instruction set design are:

1. Find keys in 32 bit set to do 128 mode transition such that existing 32 bit code can run with zero, or at worst, very simple modification. We expect good status bit communication between mode to be a significant issue.
2. Re-interpret operation of 32 instructions for operation at 128+ bits.
3. Addition of 128 bit specific inst.
4. Addition of configuration loading instruction features if necessary.
5. Addition of DPA processing instructions.

The documentation of the instructions will include binary representation, assembler phrase, and example code fragments to facilitate other members of the team to work with them.

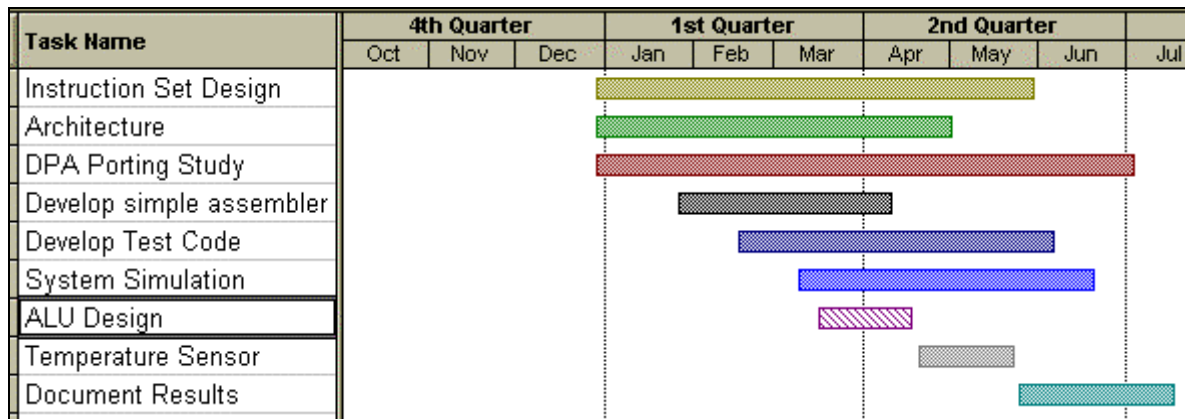


Figure 7. Research Schedule.

Architecture

The definition of the architecture will be performed primarily by Robert Heaton, Nader Bagherzadeh, a post Doctoral Associate, and a student. This activity will be performed in parallel with Instruction Set Design.

The primary form of the architecture will be as a VHDL module. Documentation will however also include a detailed top level block diagram, a top level signal definition list, and a master system timing diagram. Note that the DPA will be modeled at a very high level only during phase I.

Modeling of top level functions will be distributed among the team with one participating student performing the role of top level code project management and assembly.

DPA Porting Study

This activity will be the responsibility of one graduate student and will focus on looking at the problems of incorporating DPA logic being developed by other projects. Obviously there will be an inclination to use the architecture being developed by Obsidian in respect of BAA 97-06. However, it is not clear that that work will be sufficiently far advanced during the early part of 1998.

The result will primarily be in the form of a report, but the studies finding will obviously be taken into account by the rest of the team.

Develop Simple Assembler

A graduate student will be assigned to take this task. We will attempt to reuse an existing assembler written in C to reduce the task to about 8 weeks. The assembler will output object code in a form that can be readily used by the VHDL system simulation which may or may not be true binary.

Due to the short time available for this activity it will be limited to a small subset of the 32 bit instructions which are required for initial system simulation work.

Develop Test Code

A graduate student will be assigned to take this task and it involves taking input from the architecture, instruction set design, and assembler activities. Critical test code fragments will be developed from the available instructions and combined in a wide variety of ways to test the overall system. The activity will be supervised by Robert Heaton.

System Simulation

The graduate student responsible for supporting the architecture development will move gradually into this activity. The overall system model will be minimal, probably little more than external memory plus some method for creating interrupts. The activity will be supervised by Fadi Kurdahi.

ALU Design

Will be performed by Robert Heaton. A schematic diagram and timing diagram will be documented.

Temperature Sensor

Robert Heaton will produce schematics, spice simulation results, and a short report on the results.

Document Results

Fadi Kurdahi will supervise the assembly of the final report. This will be mainly pulling together the reports from the other activities. Robert Heaton will provide introduction, concluding remarks, and a PowerPoint format presentation based on the final report.

H. Commercial Potential.

Customer Needs

In the high end segment of the embedded processor market processing power, cost, and power are always critical for product performance. Processing power translates into faster printing, better game or simulator animations, lower cost data routing and switching. In data processing applications the 128 bit CPU offers roughly four times the throughput of a 32 bit processor.

Customer Base

The customers will be product developers who today are using 32 bit embedded processors and who will be attracted by an ARM compatible processor with much improved performance. Of course ARM, MIPS, IDT etc will not be sitting still during the three years it will take to get to market with this device. However, by putting a clear focus on high bus width devices, Obsidian plans to find a small but growing market while competing and/or collaborating with ARM and MIPS.

A typical OEM customer would be Foundry Networks. They are building low cost level 3 routers for the Internet Protocol (IP) and use the DEC StrongARM® processor which is a variant of the ARM and costs \$50 per unit. We feel that for an application like this would strongly benefit from the increased bus width, DPA instruction capability, and ARM compatibility combined with a lower price. We have prepared a preliminary data sheet to gauge the interest of a number of potential customers which is attached to this proposal.

The total market for 32 bit RISC embedded processor licensing is today approximately \$150M and growing at approximately 30% per year. The growth of the market is fueled by rapidly expanding market for intelligent products such as PDAs, low cost printers, and networking equipment.

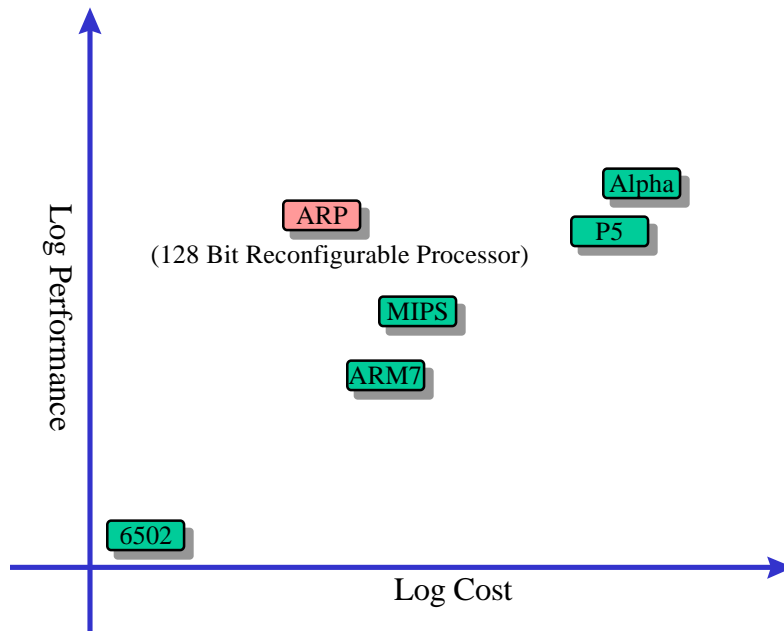


Figure 8. 128 Bit Basic Market Position.

Obsidian Technology is a startup company which will focus on development of intellectual property for licensing during the first 3 years. A number of discussions with Advanced RISC Machines Ltd., vendors of the ARM processor have occurred with respect to developing technology for ARM customers. Obsidian sees a collaborative relationship with ARM or other established players wishing to be in an ARM compatible market as key to the distribution and eventual commercial success of the proposed architecture.

Another approach to the marketing of the concept is to directly approach major high end RISC based product producers who can benefit from the additional throughput.

With the growth in interest in the IRAM approach it would also be our intention to approach a number of RAM makers to propose a joint development of an integrated processor and large SRAM or DRAM.

The background of the principle investigator Robert Heaton has a particular bearing here. He successfully negotiated the first license of the ARM RISC between Acorn Computers and VLSI Technology.

I. Principal Investigator and Senior Personnel

Robert Heaton

Robert Heaton established and managed the VLSI Design Group for Acorn Computers in 1983. He was responsible for all aspects of the group including buying the equipment and software, hiring the staff and establishing the design flow that led the effort to build the ARM, the first RISC processor chip offered for sale. Working with the team, Robert contributed the following specific architectural innovations:

1. Triple ported asymmetrical register file.
2. Semi asynchronous data processing timing flow.
3. Fast high density ALU.
4. Techniques of mixing static and half static logic.

These innovations led to a leadership position for the ARM architecture in terms of low power and silicon density which persists to this day. The ARM processor and peripheral devices were fully functional at first silicon. The ARM was implemented with an effort of just 5 man years. Robert not only led the team but was an active member of it, designing and laying out many of the key cells personally. These included the ALU, the register file, the control PLA's, the status control and address generation logic.

Since leaving Acorn during 1986, Robert has established and managed a number of engineering groups and was Senior Director of Engineering for Standard Microsystems from 1990 to 1996 where he managed a staff of 60. During this time he established a close relationship with the engineering staff in the Advanced Computer Architecture group at UC Irvine and is currently Chairman of UCI's Engineering Corporate Affiliates group.

Robert has continued to be a technical innovator with a Geometric Compression US patent #5,216,726 and a Fast Ethernet Local Area Network US patent #5,544,323. It was while working in a networking environment that Robert noticed the limitations of existing processor architectures. He established Obsidian Technology as an individual consultant during 1996 and has taken consulting contracts with Rockwell Semiconductor Systems and Standard Microsystems Corp. Additional help for this project will be drawn from UCI staff and students in the form of constancy and part time employment.

In March of this year Robert (and Obsidian Technology) was awarded an ARPA BAA 97-06 grant together with his associates at UCI.

For more information: <http://users.deltanet.com/~heaton/resume.html>

J. Consultants and Subawards.

Prof. Nader Bagherzadeh (University of California at Irvine)

Nader Bagherzadeh will be employed as a consultant to support the development of the architecture and to supervise students which work as consultants at the UCI campus. He is budgeted for 20 days. \$440 is well below his usual consulting rate.

Nader Bagherzadeh has been involved in the design of advanced microprocessor architectures for the past ten years. He was the leading designer of the first 4-issue VLIW microprocessor called VIPER. The VIPER architecture demonstrated pioneering work in the areas of processor optimization for routing operand data through bypass circuitry. Moreover, it utilized a unique pipeline architecture with a refined addressing mode that mitigated the development of future VLIW processors by industry. Later on he worked on all out of order issue Superscaler microprocessor called SDSP targeted for multimedia computation intensive applications. This study lead to several new discoveries in the areas of instruction fetching and register renaming hardware optimization.

Currently he has been combining his experience with out-of-order issue Superscaler with multi threaded architectures. In this work he is the first researcher that demonstrated the design of a register renaming scheme based on the reorder buffer and instruction window that call handle multiple threads. This work will open the way for exploiting instruction level parallelism beyond the wide issue Superscaler model. Finally, he is involved in the hardware design of a 500Mhz reorder/buffer that utilizes the True Single Phase clocking scheme.

For more information: <http://www.eng.uci.edu/comp.arch/nader/index.html>

Dr. Fadi J. Kurdahi (University of California at Irvine)

Fadi Kurdahi will be employed as a consultant to support the team's computer aided design (CAD) requirements and to supervise graduate students developing VHDL models for the project. He is budgeted for 20 days. \$440 is well below his usual consulting rate.

Fadi Kurdahi has been doing research in Design Automation and VLSI design for over 13 years. Specifically, he has researched and developed CAD tools for Architectural synthesis and estimation and has published numerous papers on the subject of linking layout information to synthesis. This becomes particularly crucial for the projected fabrication technologies that are likely to implement the proposed DRA. Recently, his work has focused on FPGAs as target implementations. His work on estimation for FPGA resulted in a suite of prediction tools that are highly accurate yet runtime efficient. Architectural synthesis techniques such as scheduling and binding are being developed to utilize these estimator and generate "first-time-correct" silicon.

Since 1987, he has been a faculty at the Department of Electrical & Computer Engineering at the University of California, Irvine, where he is currently an Associate Professor. His research interests are the areas of Computer-Aided Design of VLSI circuits, high-level synthesis, and design methodology of large scale systems. His research addresses the issues of linking High Level Synthesis to subsequent levels of design, namely logic and physical levels. He has published papers dealing with various aspects of chip and system synthesis at international conferences and journals. He has organized and participated in tutorials on high level synthesis at international conferences such as the Design Automation Conference in 1990. Prof. Kurdahi was an Associate Editor for IEEE Transactions on Circuits and Systems II in 1993-1995, and was a guest editor for a special issue of the VLSI Design Journal on "Linking Behavioral, Structural and Physical Models of Hardware". He was Organization Chairman of the Sixth International High Level Synthesis Workshop, and served on the program committees of several conferences and workshops such as: the international High Level Synthesis Workshop, the international Symposium OIL System Level Synthesis, the ACM/IEEE Physical Design Workshop, and the European Design and Test Conference. He received the Research Initiation Award from the National Science Foundation in 1989, and the ACM/SIGDA Fellowship in 1991 and 1992. Prof. Kurdahi is a member of IEEE and ACM.

For more information: <http://www.eng.uci.edu/faculty/kurdahi/fadi.html>

Prof. Nader Bagherzadeh. Consultant Letter

Dr. Fadi J. Kurdahi. Consultant Letter

NSF Budget Proposal Form 1030A

Appendix D. Supplemental: Budget Justification/Explanation Page.

- Line A1: Senior personnel. The assumption is that Robert Heaton would provide 4 days per week for this research at the daily rate of \$440. Over 6 months this amounts to \$42,240. \$440 is significantly below the \$750 per day rate Robert usually charges. Robert will be responsible for directing the project, supervising the consultants and take the leading role in developing the architecture.
- Line B1: One Post Doctoral Associate will be hired for 3 months. This is 60 working days at \$200 per day: total \$12,000.
- Line B3: Graduate Students. 5 person months @ \$15 per hour = \$12,000.
- Line B5: Secretarial - Clerical support. 10 days of general support = \$1,000.
- Line B6: Other. \$1600 for 10 days of computer systems administration support for CAD tool installation and general office automation support.
- Line E1: Domestic Travel. Estimated travel expenses for NSF conference for two persons. Also auto mileage and road tolls.
- Line G1: Materials and supplies. Cost of stationery, plotter paper, computer disks, printer heads etc. \$250.
- Line G2: Publication costs. Copying costs, scanning, postage, printing costs.
- Line G3: A total of 40 days consulting effort from Prof. Nader Bagherzadeh and Dr. Fadi Kurdahi. Please see attached letters. These consultants will be key members of the team due to their excellent background in the field and experience with managing graduate students.
- Line G4: Other. Provision of Internet services for an estimated total of 12 service months. 5 of these are for Robert Heaton, the other 7 are for the Graduate Students etc who may require at home Internet service. Hence 12 service months at \$25 per month.
Also a total of \$120 of total phone costs for the whole project. Hence total is $12 \times \$25 + \120 .
- Line I: Indirect costs. This sum is for depreciation on office equipment, computer equipment and software over 6 months, utility and accommodation costs.

References

1. [RISC] VLSI Design of the Tiny RISC Microprocessor. Arthur Abnous, Christopher Christensen, Jeffrey Gray, John Lenell, Andrew Naylor and Nader Bagherzadeh. ICS-TR-91-74. December 1991. <http://jblevins.ics.uci.edu/Dienst/UI/2.0/Describe/ncstrl.uci%2fICS-TR-91-74>
2. [DPA] Dynamically Reprogrammable Architecture: A New Approach to Adaptive Computing. Accepted proposal to ARPA BAA 97-06 by Kurdahi, Bagherzadeh, and Heaton.
3. [BRASS] Berkeley Project. <http://HTTP.CS.Berkeley.edu/projects/brass/>
4. [VIPER] Jeffrey Gray, Andrew Naylor, Arthur Abnous, and Nader Bagherzadeh, VIPER: A VLIW Integer Microprocessor.
5. [IRAM] IRAM Project. See <http://infopad.eecs.berkeley.edu/~pering/iram/project3.html>
6. [MIT] Dynamically Programmable Gate Arrays: A Step Toward Increased Computational Density (FPD '96). http://www.ai.mit.edu/projects/transit/dpga_prototype_documents.html
7. Mat Loikkanen and Nader Bagherzadeh, A Fine-Grain Multithreading Superscalar Architecture, Parallel Architectures and Compilation Techniques '96, October 1996.
8. Manu Gulati and Nader Bagherzadeh, Performance Study of a Multithreaded Superscalar Microprocessor, 2nd International Symposium on High-Performance Computer Architecture, San Jose, California, February 1996, pp 291-301.
9. Wallace, N. Dagli, and N. Bagherzadeh, Design and Implementation of a Scheduling Unit for A Superscalar Microprocessor, submitted to IEEE Transactions on VLSI Design in March 1995.
10. A. Saulsbury, et. al. "Missing the Memory Wall: The Case for Processor/Memory Integration" 23rd annual International Symposium on Computer Architecture, June 1996.
11. [VLSI] Acorn RISC Machine. VLSI Design Conference. Robert Heaton, November 1985.

This project will be conducted in association with the staff and students of the Advanced Computer Architecture Group of the University of California at Irvine.

K. Equipment, Instrumentation, Computers, and Facilities

The facilities of the UCI Computer Labs will be used by the consultants as a part of there contract. Other work will be performed in home offices and are thus available. The equipment needs for the project are modest:

1. 4 Personal Computers with Word Processing, and VHDL simulation. These will be owned by Obsidian Technology.
2. Various printers and other office automation equipment owned by Obsidian.
3. OrCad, Mentor QuickVHDL, and Tanner design tools owned by Obsidian.

L. Current and Pending Support of the Principal Investigator and Senior Personnel.

The PI and two consultants have recently been awarded a DARPA grant BAA 97-06 with the title "Dynamically Reconfigurable Architecture: A New Approach to Adaptive Computing". Funding for Obsidian (as a subcontractor to UCI) is expected to start during August 1997. The total Obsidian funding over three years is \$300,000 and represents a time commitment from the PI of 3.7 person-months per year.

M. Equivalent or Overlapping Proposals to Other Federal Agencies

None. This project may use the work of the DARPA project in L above. However we feel there is no overlap in the proposals because this proposal focuses on the use rather than the development of DPA cells and design tools.

O. Prior SBIR Awards

None.

Example APR128 Data Sheet