Architecture and Evaluation of An Asynchronous Array of Simple Processors

Zhiyi Yu, Michael J. Meeuwsen, Ryan W. Apperson, Omar Sattari, Michael A. Lai,
Jeremy W. Webb, Eric W. Work, Tinoosh Mohsenin, Bevan M. Baas

Contact: zhyyu@ucdavis.edu, bbaas@ucdavis.edu

Abstract— **This paper presents the architecture of an Asynchronous Array of simple Processors (AsAP), and evaluates its key architectural features as well as its performance and energy efficiency. The AsAP processor calculates DSP applications with high energy-efficiency, is capable of high-performance, is easily scalable, and is well-suited to future fabrication technologies. It is composed of a 2-D array of simple single-issue programmable processors interconnected by a reconfigurable mesh network. Processors are designed to capture the kernels of many DSP algorithms with very little additional overhead. Each processor contains its own tunable and haltable clock oscillator, and processors operate completely asynchronously with respect to each other in a globally asynchronous locally synchronous (GALS) fashion.**

A 6×6 AsAP array has been designed and fabricated in a 0.18 μm CMOS technology. Each processor occupies 0.66 mm^2, is fully functional at a clock rate of 520–540 MHz at 1.8 V, and dissipates an average of 35 mW per processor at 520 MHz under typical conditions while executing applications such as a JPEG encoder core and a complete IEEE 802.11a/g wireless LAN baseband transmitter. Most processors operate at over 600 MHz at 2.0 V. Processors dissipate 2.4 mW at 116 MHz and 0.9 V. A single AsAP processor occupies 4% or less area than a single processing element in other multi-processor chips. Compared to several RISC processors (single issue MIPS and ARM), AsAP achieves performance 27–275 times greater, energy efficiency 96–215 times greater, while using far less area. Compared to the TI C62x high-end DSP processor, AsAP achieves performance 0.8–9.6 times greater, energy efficiency 10–75 times greater, with an area 7–19 times smaller. Compared to ASIC implementations, AsAP achieves performance within a factor of 2–5, energy efficiency within a factor of 3–50, with area within a factor of 2.5–3. These data are for varying numbers of AsAP processors per benchmark.

I. Introduction

Applications that require the computation of complex DSP workloads are becoming increasingly commonplace. These applications often comprise multiple DSP tasks and are found in applications such as: wired and wireless communications, multimedia, sensor signal processing, and medical/biological processing. Many are embedded and are strongly energy-constrained. In addition, many of these workloads require very high throughputs and often dissipate a significant portion of the system power budget and are therefore of considerable interest.

Increasing clock frequencies and an increasing number of circuits per chip has resulted in modern chip performance being limited by power dissipation rather than circuit constraints. This implies a new era of high-performance design that must now focus on energy-efficient implementations [1]. Future fabrication technologies are expected to have large variations in devices and wires, and "long" wires are expected to significantly reduce maximum clock rates. Therefore, architectures that enable the elimination of long high-speed wires will likely be easier to design and may operate at higher clock rates [2].

The Asynchronous Array of simple Processors (AsAP) computes the aforementioned complex DSP application workloads with high performance and high energy-efficiency, and is well suited for future technologies. The AsAP system comprises a 2-D array of simple programmable processors interconnected by a reconfigurable mesh network. Processors are each clocked by fully independent haltable oscillators in a globally asynchronous locally synchronous (GALS) [3] fashion. Several of AsAP's key features distinguish it from other broadly similar work:

- **A chip multiprocessor architecture** achieves high performance through parallel computation. Many DSP applications are composed of a collection of cascaded DSP tasks, so an architecture that allows the parallel computation of independent tasks will likely be more efficient.
- **Small memories and simple single-issue architecture** for each processor achieves high energy efficiency. Since large memories—which are normally used in modern processors [4], [5]—dissipate significant energy and require larger delays per memory transaction, architectures that minimize the need for memory and keep data near or within processing elements are likely to be more efficient. Along with reduced memory sizes, the datapath and control logic complexity of AsAP are also reduced.
- **GALS clocking style** is suitable for future fabrication technologies and can achieve high energy efficiency due to the fact that global clock circuits have become increasingly difficult to design and they consume significant power.
- **Nearest neighbor communication** is used to avoid global wires to make it suitable for future fabrication technologies, due to the fact that global chip wires will dramatically limit performance if not properly addressed since their delay is roughly constant when scaled [2].

A prototype 6×6 AsAP chip has been implemented in 0.18 μm CMOS and is fully functional [6]. In this paper, we discuss AsAP's architectural design and investigate how the key features affect system results. In addition, we present a thorough evaluation of its performance and energy efficiency

TABLE I

ASAP 32-BIT INSTRUCTION TYPES AND FIELDS

Instruction Type	6 bits	8 bits	8 bits	8 bits	2 bits
General	opcode	dest	src1	src2	NOPs
Immediate	opcode	dest	immediate		NOPs
Branch	opcode	–	–	target	NOPs

TABLE II

CLASSES OF THE 54 SUPPORTED INSTRUCTIONS

Instruction class	Number of instructions
Addition	7
Subtraction	7
Logic	11
Shift	4
Multiply	2
Multiply-accumulate	6
Branch	13
Miscellaneous	4

for several DSP applications.

II. THE ASAP PROCESSOR SYSTEM

A. Architecture of the AsAP processor

The AsAP array consists of a large number of simple uniform processing elements operating asynchronously with respect to each other and connected through a reconfigurable network. The processors are optimized to efficiently compute DSP algorithms individually as well as in conjunction with neighboring processors. Figure 1 contains diagrams of the fabricated processing array and a single AsAP processor.

Each AsAP processor is a simple single-issue processor with a 64-word 32-bit instruction memory (IMEM), a 128-word 16-bit data memory (DMEM), a dynamic configuration memory (DCMEM), a 16×16-bit multiplier with a 40-bit accumulator, a 16-bit ALU, and four address generators. It utilizes a memory-to-memory architecture with no register file. No support is provided for branch prediction, out of order execution, or speculative operation. During the design phase, hardware was added only when it significantly increased performance and/or energy-efficiency for our benchmarks. A nine stage pipeline is implemented as shown in Fig. 1. All control signals are generated in the instruction decode stage, and pipelined appropriately. Interlocks are not implemented in hardware, so all code is scheduled prior to execution by the compiler.

1) Instruction set: AsAP supports 54 32-bit instructions with three broad instruction formats. A summary of the 54 instructions is given in Tables I and II. *General* instructions select two operands from memories, accumulator, FIFOs, and three ALU bypass routes; and they select one destination from memories, accumulator and output ports. *Immediate* instructions receive input from a 16-bit immediate field. *Branch* instructions include a number of conditional and unconditional branch functions. Two bits in each instruction define how many NOP operations (from 0 to 3) should follow after instruction processing, which allows inserting NOPs to avoid pipeline hazards without requiring additional NOP instructions.

TABLE III

DATA FETCH ADDRESSING MODES

Addressing mode	Example	Meaning
Direct	Move Obuf DMEM 0	Obuf ← DMEM[0]
Address pointer	Move Obuf aptr0	Obuf ← DMEM[DCMEM]
Address generator	Move Obuf ag0	Obuf ← DMEM[generator]
Short immediate	Add Obuf #3 #3	Obuf ← 3+3
Long immediate	Move Obuf #256	Obuf ← 256
DCMEM	Move Obuf DCMEM 0	Obuf ← DCMEM[0]
Bypassing	Move Obuf regbp1	Obuf ← first bypass
FIFOs	Move Obuf Ibuf0	Obuf ← FIFO 0
ACC	Move Obuf Acc	Obuf ← ACC[15:0]

Other than a bit-reverse instruction and a bit-reverse mode in the address generators, no algorithm-specific instructions or hardware are implemented. While single-purpose hardware can greatly speed computation for specific algorithms, it can prove detrimental to the performance of a complex multi-algorithmic system and limits performance for future presently-unknown algorithms—which is one of the key domains for programmable processors.

2) Data addressing: AsAP processors fetch data at pipeline stage *Mem Read*, using the addressing modes listed in Table III. Three methods are supported to address data memory. *Direct* memory addressing uses immediate data as the address to access static memory locations; four *address pointers* access memory according to the value in special registers located in DCMEM; and four *address generators* provide automatic addressing with special-purpose hardware to accelerate many tasks. In addition to the data memory, AsAP processors can also fetch data from another 6 locations: 1) short immediate data (6 bits) can be used in dual-source instructions, 2) long immediate data (16 bits) can be used in the move immediate instruction, 3) the DCMEM's configuration information can be read or written by instructions, 4) three bypass paths from the ALU and MAC units can be used as sources to accelerate execution, 5) the two processor input FIFOs are available as general instruction sources, and 6) the lowest 16 bits of the accumulator register can also be a source for execution.

Figure 2 shows the logic diagram for one address generator. Each address generator contains a *count* register which is used as the memory pointer, and several inputs define how to change its value after each memory access. *Start_addr* defines the start address of the *count* register. When the counter is enabled (*enable* = 1), it will be increased or decreased (determined by *direction*) by the amount of the value *stride*. The *count* register is reloaded to the start address when it reaches the end address (*end_addr*). The other control signals are primarily used to accelerate FFTs. Each address generator occupies about 3700 μm^2 in a 0.18μm technology, and the four address generators occupy only 2% of the processor's area.

3) Completely independent clocking and circuits for crossing asynchronous clock domains: Each processor has its own digitally programmable clock oscillator which occupies only about 0.5% of the processor's area. There are no phase-locked loops (PLLs), delay-locked loops (DLLs), or global frequency or phase-related signals, and the system is fully GALS. While impressive low clock skew designs have been

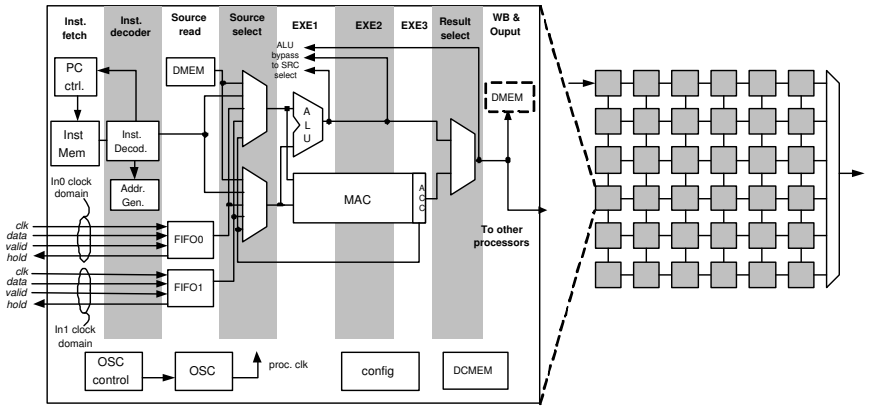

Fig. 1. Block diagram of an AsAP processor and the 6×6 chip. Vertical gray bars indicate the nine pipeline stages.

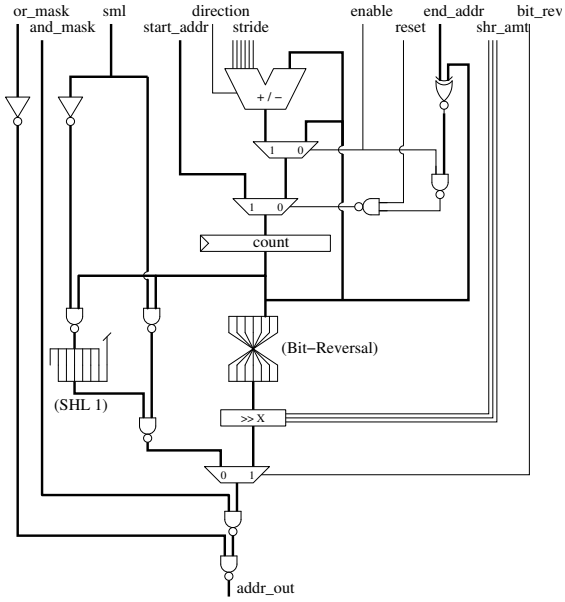

Fig. 2. Address generator; thin lines represent one-bit wires, and thick lines represent seven-bit wires.

achieved at multi-GHz clock rates, the effort expended in clock tree management and layout is considerable [7]. Placing a clock oscillator inside each processor reduces the size of the clock tree circuit to a fraction of a mm^2—the size of a processing element. Large systems can be made with arrays of processing elements with no change whatsoever to clock trees (that are wholly contained within processing elements), simplifying overall design complexity and scalability.

The reliable transfer of data across unrelated asynchronous clock domains is accomplished by dual-clock FIFOs [8]. The FIFO's write clock and data are supplied in a source-synchronous fashion by the upstream processor and its read clock is supplied by the downstream processor—which is the host for the dual-clock FIFO in AsAP.

Special clock control circuits enable processing elements to power down completely—dissipating leakage power only—if no work is available for nine clock cycles. The local oscillator is fully restored to full speed in less than one cycle after work again becomes available.

4) Reconfigurable 2-D mesh network: Processors connect via a configurable 2-dimensional mesh. To maintain link communication at full clock rates, inter-processor connections are made to nearest-neighbor processors only. A number of architectures including wavefront [9], RAW [10], and TRIPS [11], have specifically addressed this concern and have demonstrated the advantages of a tile-based architecture. AsAP's nearest neighbor connections result in no high-speed wires with a length greater than the linear dimension of a processing element. The inter-processor delay decreases with advancing fabrication technologies and allows clock rates to scale upward. Longer distance data transfers in AsAP are handled by routing through intermediary processors or by "folding" the application's data flow graph so that communicating processing elements are placed adjacent or near each other—for example, the *Pilot Insert* processor and the first *G.I. Wind.* processor in Fig. 5b.

Each AsAP processor has two asynchronous input data ports and can connect each port to any of its four nearest neighboring processors. Because changing active clock signals can cause runt clock pulses, a processor may change its input connection only during times when both input clocks are guaranteed to both be low—which is normally only during power-up. On the other hand, output port connections can be changed among any combination of the four neighboring processors at any time through software.

B. AsAP implementation

The first generation AsAP 6×6 processor array has been implemented using TSMC 0.18 μm CMOS technology [6]. The left part of Fig. 3 shows the die micrograph. A standard cell based design flow was used from verilog source code. All circuits were synthesized, except the programmable oscillator. A single processor tile and the entire chip were placed and routed by CAD tools.

The right part of Figure 3 shows the test environment for the AsAP prototype including a printed circuit board hosting an AsAP processor and a supporting FPGA board to interface between AsAP and a host PC. AsAP's SPI-style serial

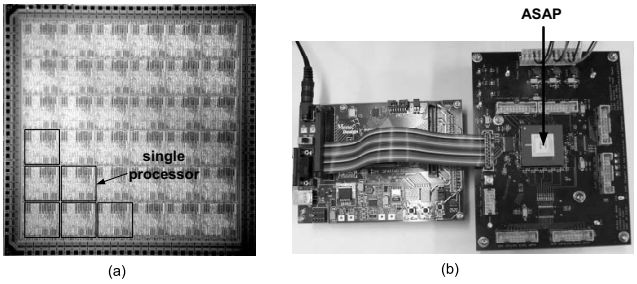

Fig. 3. Micrograph of the (a) 6×6 AsAP chip and (b) its test environment

```
void main() {
  int i;
  int a, b, c;
  sat int d;          /* saturating integer */
  while(1) {          /* loop */
    a = Ibuf0;        /* read value from FIFO 0 */
    b = Ibuf1;        /* read value from FIFO 1 */
    c = (a + b) >> 1;/* AddHigh instruction */
    for (i = 0; i < 10; i++) {
      d = c + i;      /* saturating instruction */
      OBuf = d;       /* OBuf is proc.'s output */
    }
  }
}
```

Fig. 4. An example C language program for an AsAP processor

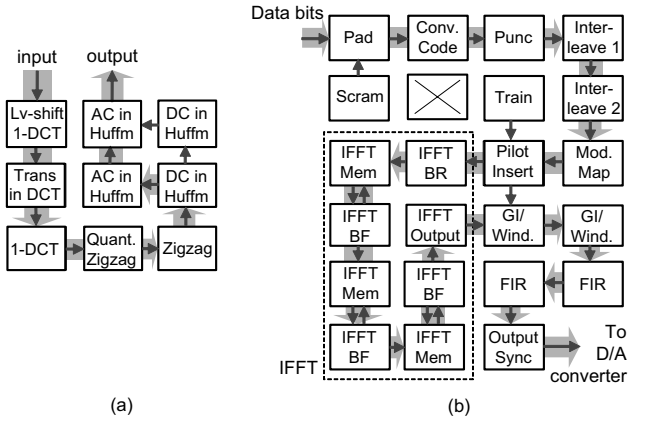

Fig. 5. Block diagram of the (a) 9-processor JPEG encoder and (b) 22-processor 802.11a/g implementation. Thin arrows show all paths and wide arrows show the primary data flow. The processor marked with an "×" is unused and powered down.

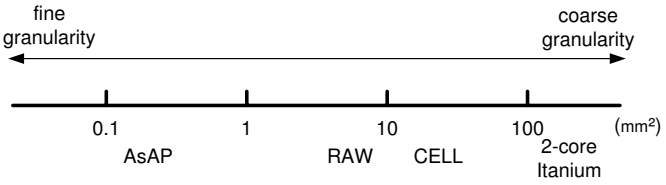

Fig. 6. Size of a single processing element in several chip multi-processor systems. Data are scaled to 0.13 μm CMOS technology.

port receives configuration information and programs for each processor.

C. Software

Programming the AsAP processor presents significant challenges. Programming involves taking advantage of all levels of parallelism easily available to simplify the coding of small kernels, including task-level parallelism, data-level parallelism, and address-data parallelism. Partly due to the natural partitioning of applications by task-level parallelism, we have found the task less challenging than first expected. This is supported by data in Table IV showing the memory requirements of common DSP tasks.

A high level language (which we call called AsAP-C) and its corresponding compiler were developed to generate code for each individual AsAP processor. AsAP-C contains most standard C language functions such as arithmetic calculations (addition, subtraction, multiplication, etc.), logic calculations (AND, OR, NOT, etc.), and control functions (while loops, for loops, etc.). A saturating integer type is defined to support DSP integer calculations which are commonly used in high level DSP languages [12]. Additionally, the language contains several functions specific for AsAP such as FIFO reads and direct inter-processor communication. Both inputs and outputs are mapped into the language through the reserved variable names: `Ibuf0`, `Ibuf1`, and `Obuf`. Figure 4 shows one example of AsAP-C program which fetches data from two FIFOs and sends its result to the processor's output port.

The job of programming processors also includes the mapping of processor graphs to the 2-D planar array. While this is normally done at compile time, an area of current work is tools for the automatic mapping of graphs to accommodate rapid programming and to recover from hardware faults and extreme variations in circuits, environment, and workload.

D. Task and application implementations

In addition to the tasks listed in Table IV, we have completed the implementation and further analysis of several complex applications, including a JPEG encoder and an IEEE 802.11a/g wireless LAN baseband transmitter as shown in Fig. 5—both are fully functional on the fabricated chip. The JPEG encoder principally consists of five sub-tasks: level shift, 8×8 Discrete Cosine Transform (DCT), quantization, zig-zag scanning, and Huffman encoding. The fully-compliant 802.11a/g transmitter implementation operates over all 8 data rates, includes additional upsampling and filtering functions not specified by the standard [13], and sustains transmissions at 30% of the full 54 Mb/s rate at a clock speed of 300 MHz on 22 processors [14].

III. Analysis of the Key Features

One of the most important variables in chip multiprocessor architectures is the level of granularity of each processing element. A wide range of granularities are possible as shown in Fig. 6 [6], [10], [15], [16]. The coarse grain 2-core Itanium [16] contains large wide-issue processors each close to 300 mm^2 in 90 nm technology, while the fine grain AsAP contains single-issue processors less than 1 mm^2 in 0.18 μm

TABLE IV

REQUIRED INSTRUCTION MEMORY AND DATA MEMORY SIZES FOR VARIOUS DSP TASKS ON A SIMPLE SINGLE-ISSUE PROCESSOR

Task	IMem Size (words)	DMem Size (words)
16-tap FIR filter	6	33
Level-shifting for JPEG	8	1
8-point DCT	40	16
8×8 2-D DCT	154	72
Quantization for 64 elements	7	66
Zig-zag re-ordering for JPEG	68	64
Huffman encoding for JPEG	203	334
Scrambling for 802.11a/g	31	17
Padding OFDM bitstream	49	25
Convolutional coding ($k = 7$)	29	14
Interleaving 1 for 802.11a/g	35	30
Interleaving 2 for 802.11a/g	51	31
Modulation for BPSK, QPSK, 16QAM, 64QAM	53	33
Pilot insertion for OFDM	47	68
Training symbol generation for 802.11a/g	31	76
64-pt complex FFT	97	192
Guard interval insertion for OFDM	44	74
2× upsampling + 21-tap Nyquist FIR filter	40	128
Bubble sort	20	1
N merge sort	50	N
Square root	62	15
Exponential	108	32

technology. Size differences of factors of tens and hundreds make strong impacts on system behavior.

Most chip multiprocessors target a broad range of applications, and each processor in such systems normally contains powerful computational resources—such as large memories, wide issue processors [16], and powerful inter-processor communication [10]—to support widely varying requirements. Extra computational resources can enable systems to provide high performance to a diverse set of applications, but they reduce energy efficiency for tasks that can not make use of those specialized resources. Most DSP applications AsAP targets are made up of computationally intensive tasks with very small instruction and data kernels, which makes it possible to use extremely simple computational resources—small memory, simple single issue datapath, and nearest neighbor communication—to achieve high energy efficiency while maintaining high performance.

In this section, we analyze these key features of the AsAP processor which justify its fine grain architecture. We also briefly analyze AsAP's GALS clocking style.

A. Small memory

A clear trend among all types of programmable processors is not only an increasing amount of on-chip memory, but also an increasing percentage of die area used for memory. For example, the TI C64x [4] and third generation Itanium processor [5] use approximately 75% and 70% area for memory respectively. Since large memories dissipate more energy and require larger delays per transaction, we seek architectures that minimize the need for memory and keep data near or within processing elements.

1) Inherent small memory requirement for DSP applications: A notable characteristic of the targeted DSP tasks is that many have very limited memory requirements compared to general-purpose tasks. The level of *required* memory must be differentiated from the amount of memory that can be used or is typically used to calculate these kernels. For example, an N-tap filter may be programmed using a vast amount of memory though the base kernel requires only $2N$ data words. Table IV lists the actual amounts of instruction and data memory required for 22 common DSP tasks and shows the very small required memory sizes compared to memories commonly available in modern processors. This analysis assumes a simple single-issue processor like AsAP. Although programs were hand written in assembly code, little effort was placed on optimizing them such as scheduling instructions for the pipeline or using forwarding paths.

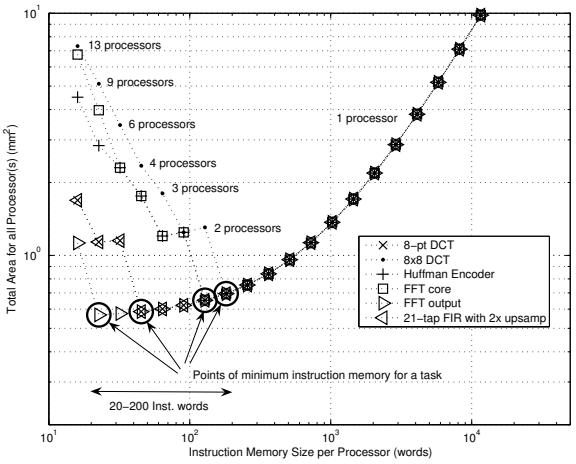

Fig. 7. Total area required for representative tasks mapped onto one or multiple 0.18 μm CMOS simple processors, as a function of the size of the instruction memory per processor. Minimum total area is achieved with approximately 20–200 instructions per processor.

2) Finding the optimal memory size for DSP applications: Once the amount of required instruction and data memory is known, it is worthwhile to consider what size of memory per processor is optimal in terms of total processing element area. We begin our analysis with several assumptions: 1) the non-memory processor size is 0.55 mm^2 in 0.18 μm CMOS and is not a function of memory size, 2) memory area scales linearly with capacity and the area is 400 $\mu\mathrm{m}^2$ for a 16-bit word and 800 $\mu\mathrm{m}^2$ for a 32-bit word, 3) a fixed partitioning overhead is added each time a task is split onto multiple processors—this overhead is estimated per task and varies from 2–8 instructions and from 0–30% of the total space, and 4) additional processors used only for routing data may be needed for designs using a large number of processors, but are neglected. Figures 7 and 8 show the total circuit areas for several representative tasks listed in Table IV, while varying the instruction memory and data memory sizes respectively.

These analyses show that processors with memories of a few hundred words will likely produce highly energy efficient systems due to their low overall memory power and their very short intra-processor wires. On the negative side, processors with very small memories that require parallelization of tasks across processors may require greater communication energy

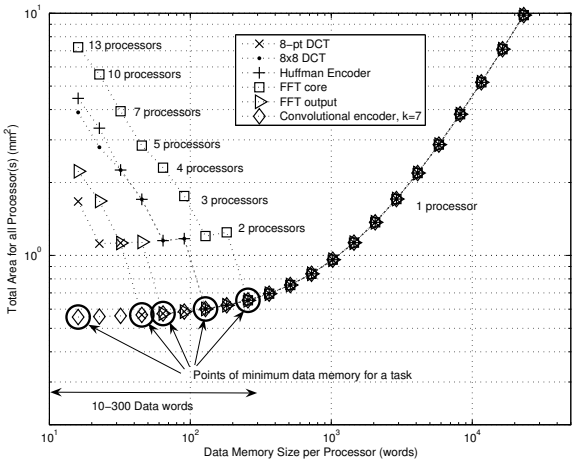

Fig. 8. Analysis similar to that shown in Fig. 7 but for data memory. The minimum total area is achieved with approximately 10–300 data words per processor.

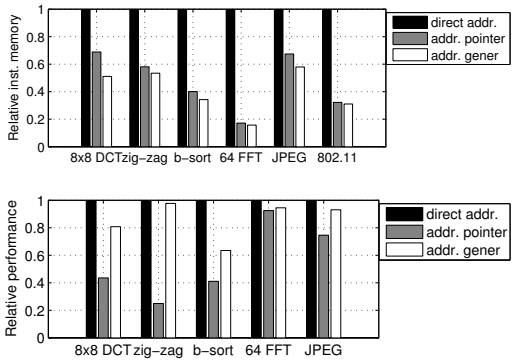

Fig. 9. Comparison of relative instruction memory requirements and relative performance for three different addressing modes. Comparisons are made against the *direct address* case which uses straight line code with pre-calculated addresses only.

and present significant programming challenges.

3) Several architectural features help reduce memory requirement: In addition to the inherent small instruction memory requirement of DSP applications, address generators help reduce the required instruction memory for applications since they can handle many complex addressing functions without any additional instructions. The upper figure of Fig. 9 shows the estimated relative instruction cost for a system using three addressing modes to fulfill the same functions. Compared to systems primarily using direct memory addressing and address pointers, the system containing address generators reduces the number of required instructions by 60% and 13% respectively. Also, using address generators can increase system performance. As shown in the lower figure of Fig. 9, it comes within 15% of the performance of a system using direct addressing with pre-calculated addresses, and approximately 2 times higher performance compared to a system using address pointers alone.

The embedded NOP instruction field described in Sec. II-A.1 also helps reduce instruction memory requirements since it

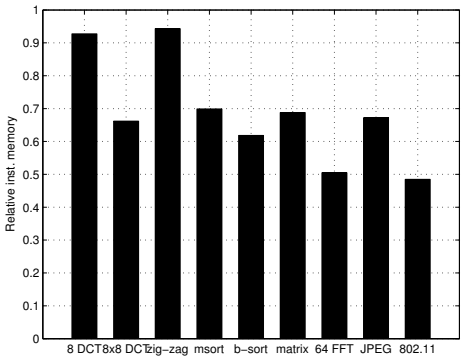

Fig. 10. Relative instruction memory reductions by using a 2-bit embedded NOP field in each instruction

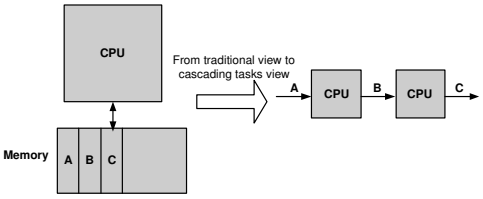

Fig. 11. Traditional large memory versus small memory cascading tasks views

dramatically reduces the number of explicit NOP instructions. Figure 10 shows that instruction memory requirements can be reduced by approximately 30% for 9 applications.

In addition to the inherent small data memory requirements of DSP applications, task cascading also helps to reduce the required data memory size. As shown in Fig. 11, a system with many processors can use separate processors to compute individual tasks in an application, and the intermediate data can be streamed between processors instead of buffering them in a large memory.

B. Datapath—wide issue vs. single issue

The datapath, or execution unit, plays a key role in processor computation, and also occupies a considerable amount of chip area. Uniprocessor systems are shifting from single issue architectures to wide issue architectures in which multiple execution units are available to enhance system performance. For chip multiprocessor systems, there remains a question about the trade-off between using many small single-issue processors, versus larger but fewer wide-issue processors.

A large wide-issue processor has a centralized controller, contains more complex wiring and control logic, and its area and power consumption increase faster than linearly along with the number of execution units. One model of area and power for processors with different issues derived by J. Oliver et al. [17] shows using wide-issue processors consumes significantly more area and power than using multiple single-issue processors. Their work shows a single 32-issue processor occupies more than 2 times the area and dissipates approximately 3 times the power of 32 single-issue processors.

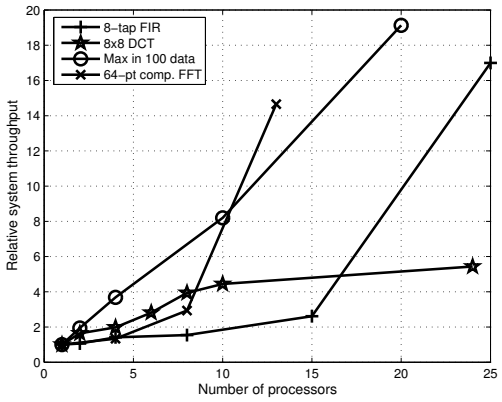

Fig. 12. Increase in system throughput with increasing number of processors

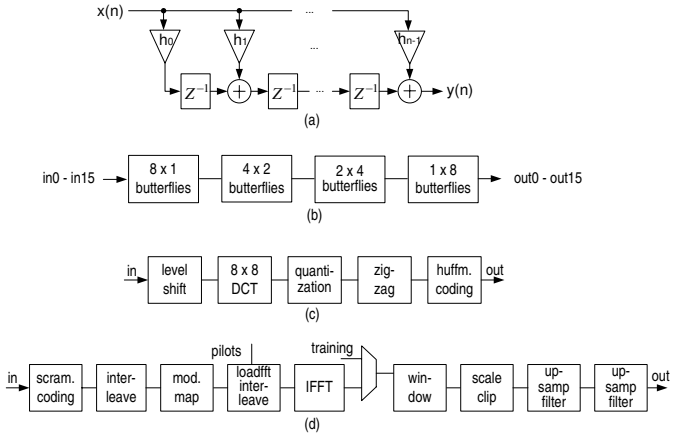

Fig. 13. Common DSP tasks and applications exhibiting their abilities to be linearly pipelined: (a) transpose form FIR filter, (b) 16-point radix-2 FFT, (c) JPEG encoder, and (d) IEEE 802.11a/g wireless LAN transmitter

However, chip multiprocessor systems composed of single-issue processors will not always have higher area and energy efficiency—much depends on specific applications. Wide-issue processors work well when instructions fetched during the same cycle are highly independent and can take full advantage of functional unit parallelism, but this is not always the case. Multiple single-issue processors such as AsAP are less efficient if the application is not easy to partition, but it can perform particularly well on many DSP applications since they are often made up of complex components exhibiting task level parallelism so that tasks are easily spread across multiple processors. Large numbers of simple processors also introduce extra inter-processor communication overhead, which we discuss further in Sec. III-C.

Figure 12 shows how throughput scales for four single tasks relative to the throughput of a single processor. Programs were written in assembly by hand but are lightly optimized and unscheduled. The memory requirement for the 8×8 DCT and 64-pt complex FFT exceeds the available memory of a single AsAP processor, so data points using one processor are estimated assuming one single processor had a large enough memory. An analysis of scaling results of a 16-tap FIR filter implemented in 85 different designs using from 1–52 processors shows a factor of 9 variation in throughput per processor over this space [18].

When all processors have a balanced computational load with little communication overhead, the system throughput increases close to linearly with the number of processors, such as for the task of finding the maximum value of a data set (*Max in 100 data* in Fig. 12). Clearly, applications that are difficult to parallelize show far less scalability at some point. For example, the performance of the 8×8 DCT increases well up to 10 processors where 4.4 times higher performance is achieved, but after that, little improvement is seen and only 5.4 times higher performance is seen using 24 processors. However, there is significant improvement in the FIR filter and FFT after a certain number of processors is reached. The reason for this is because increasing the number of processors in these applications avoids extra computation in some cases. For example, the FFT avoids the calculation of data and coefficient addresses when each processor is dedicated to one stage of the FFT computation. On average, 10 processor and 20 processor systems achieve more than 5 times and 10 times higher performance compared to a single processor system, respectively.

C. Nearest neighbor communication

Currently, most chip multiprocessors target broad general purpose applications and use complex inter-processor communication strategies [10], [19], [20], [21], [22]. For example, RAW [10] uses a separate complete processor to provide powerful static routing and dynamic routing functions, Blue-Gene/L [21] uses a torus network and a collective network to handle inter-processor communication, and Niagara [22] uses a crossbar to connect 8 cores and memories. These methods provide flexible communication abilities, but consume a significant portion of the area and power in communication circuits.

The DSP applications which AsAP targets have specific regular features and make it possible to use a simple nearest neighbor communication scheme to achieve high area and energy efficiency, without a large performance loss. As can be seen from several popular industry-standard DSP benchmarks such as TI [23], BDTI [24], and EMBC [25], the most common tasks include FIR and IIR filtering, vector operations, the Fast Fourier Transform (FFT), and various control and data manipulation functions. These tasks can normally be *linearly pipelined*, as shown in the upper two examples in Fig. 13, and the result from one stage can be sent directly to the next stage without complex global communication. Complete applications containing multiple DSP tasks also have similar features, as two examples shown in Fig. 13(c) and (d) for the JPEG encoder and the 802.11a/g baseband transmitter. All these examples can be handled efficiently by nearest neighbor inter-processor communication.

Nearest neighbor communication simplifies the inter-processor circuitry and two dual-clock FIFOs present the major cost in this case, which results in low area and high energy efficiencies. Figure 14 compares AsAP to four other chip

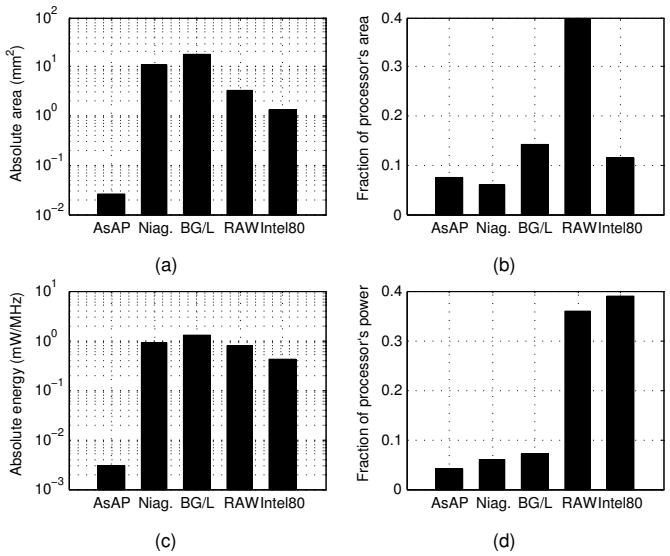

Fig. 14. Comparison of communication circuit area and power dissipation for five chip multiprocessors: (a) absolute area; (b) fraction of processor's circuit area used for communication; (c) absolute energy dissipation; and (d) fraction of processor's power dissipation due to communication circuits. Values are scaled to 0.13 μm CMOS technology.

multiprocessors (Niagara [22], BlueGene/L [21], RAW [20], and Intel 80-core [26]). The communication circuitry in the AsAP processor occupies less than 0.08 mm^2 in 0.18 μm CMOS, which is approximately 8% of the processor area, and is more than 50 times smaller than the others. Under the worst case conditions when maximizing possible communication, the communication circuity in the AsAP processor consumes around 4 mW at 475 MHz, which is about 4% of the processor power, and the energy efficiency is more than 100 times higher than the others.

D. GALS

The GALS clocking style simplifies the clock tree design and provides the opportunity for a joint clock/voltage scaling method to achieve very high energy efficiency. However, at the same time, it introduces an extra performance penalty since it requires extra circuitry to handle asynchronous boundaries which introduce additional latency. It has been shown that the performance penalty from a GALS chip multiprocessor architecture like AsAP can be highly reduced, due to its localized computation and less frequent communication loops. Simulation results show the performance penalty of the AsAP processor is less than 1% compared to the corresponding synchronous system [27].

IV. Evaluation of the AsAP processor

This section provides a detailed evaluation and discussion of the AsAP processor including performance, area, and power consumption.

Each processor occupies 0.66 mm^2 and the 6×6 array occupies 32.1 mm^2 including pads. Due to its small memories and simple architecture, each AsAP processor's area is divided as follows: 8% for communication circuitry, 26% for memory circuitry, and a favorable 66% for the remaining core.

TABLE V
Estimates for a 13 mm × 13 mm AsAP array implemented in various semiconductor technologies

CMOS Tech (nm)	Processor Size (mm^2)	Num Procs per Chip	Clock Freq (GHz)	Peak System Processing (Tera-Op)
180	0.66	256	0.51	0.14
130	0.34	500	0.66	0.33
90	0.16	1050	1.02	1.07
45	0.04	4200	2.04	8.57

Processors operate at 520–540 MHz under typical conditions. The average power consumption for each processor is 35 mW when processors are executing applications such as a JPEG encoder or an 802.11a/g baseband transmitter, and they consume 94 mW when 100% active at 520 MHz. At a supply voltage of 2.0 V, most processors operate at clock frequencies over 600 MHz.

A. High speed, small area, and high peak performance

Small memories and simple processing elements enable high clock frequencies and high system performance. The AsAP processor operates at frequencies among the highest possible for a digital system designed using a particular design approach and fabrication technology. The clock frequency information listed in Table VI supports this assertion.

AsAP is also highly area efficient. AsAP has a processing element density about 23–100 times greater than that of other broadly-similar projects [6], and thus each AsAP processor occupies 4% or less area compared to other reported processing elements.

High clock speeds and small area result in a high peak performance density with a fixed chip size. With advancing semiconductor fabrication technologies, the number of processors will increase with the square of the scaling factor and clock rates will increase approximately linearly—resulting in a total peak system throughput that increases with the *cube* of the technology scaling factor. Table V summarizes area and performance estimates for several technologies with the corresponding peak performance. It shows that in 90 nm technology, an AsAP array can achieve 1 Tera-op/sec with a 13 mm×13 mm chip. Real applications would unlikely be able to sustain this peak rate, but tremendous throughputs are nonetheless expected.

B. High performance and low power consumption for DSP applications

Table VI lists area, performance, and power data for a number of general-purpose, programmable DSP, and ASIC processors for which we could obtain data. We choose the TI C62x as the reference programmable DSP processor since it belongs to the TI VLIW C6000 series which is TI's highest performance series. The enhanced TI C64x VLIW DSP processor [4] is also in the C6000 series and has an architecture similar to the C62x, but it contains substantial

TABLE VI

AREA, PERFORMANCE AND POWER COMPARISON OF VARIOUS PROCESSORS FOR SEVERAL KEY DSP KERNELS AND APPLICATIONS; ALL DATA ARE SCALED TO 0.18 μm TECHNOLOGY ASSUMING A $1/s^2$ REDUCTION IN AREA, A FACTOR OF s INCREASE IN SPEED, AND A $1/s^2$ REDUCTION IN POWER CONSUMPTION. THE AREA IS THE CHIP CORE AREA WHEN AVAILABLE.

Benchmark	Processor	Processor style	Scaled area (mm^2)	Scaled clock freq. (MHz)	Clock cycles	Scaled execution time (ns)	Scaled power (mW)	Scaled energy (nJ)
40-tap FIR filter	**AsAP**	**Array (8 proc.)**	**5.28**	**510**	**10**	**20**	**730**	**15**
	MIPS VR5000 [23], [28]	RISC processor.	N/A	347	430	1239	2600	3222
	TI C62x [23], [28]	VLIW DSP	> 100	216	20	92	3200	296
	PipeRench [29]	Parallel processor	55	120	2.87	24	1873	45
8x8 DCT	**AsAP**	**Array (8 proc.)**	**5.28**	**510**	**254**	**498**	**390**	**194**
	NMIPS [30]	RISC processor	N/A	78	10772	137000	177	24400
	CDCT6 [30]	Enhanced RISC	N/A	178	3208	18000	104	1950
	TI C62x [23], [28]	VLIW DSP	> 100	216	208	963	3200	3078
	DCT processor [31]	ASIC	1.72	555	64	115	520	60
Radix-2 complex 64-pt FFT	**AsAP**	**Array (13 proc.)**	**8.6**	**510**	**845**	**1657**	**730**	**1209**
	MIPS VR5000 [28]	RISC proc.	N/A	347	15480	44610	2600	115988
	TI C62x [23], [28]	VLIW DSP	> 100	216	860	3981	3200	12739
	FFT processor [32]	ASIC	3.5 (core)	27	23	852	43	37
JPEG encoder (8x8 block)	**AsAP**	**Array (9 proc.)**	**5.94**	**300**	**1443**	**4810**	**224**	**1077**
	ARM [33]	RISC processor	N/A	50	6372	127440	N/A	N/A
	TI C62x [33], [28]	VLIW DSP	> 100	216	840	3900	3200	12400
	ARM+ASIC [33]	ASIC + RISC	N/A	50	1023	20460	N/A	N/A
802.11a/g transmitter (1 symbol)	**AsAP**	**Array (22 proc.)**	**14.52**	**300**	**4000**	**13200**	**407**	**5372**
	TI C62x [34] [28]	VLIW DSP	> 100	216	27200	126800	3200	405760
	Atheros [35]	ASIC	4.8 (core)	N/A	N/A	4000	24.2	96.8

circuit level optimizations that achieve more than 4 times higher performance with less than half the power consumption compared to the C62x. We feel the C62x is a fair comparison with the first version AsAP processor and thus a better comparison at the architecture level, without tainting from circuit level optimizations.

In support of our assertion that the AsAP prototype has significant room for improvement, we note that measurements show approximately 2/3 of AsAP's power is dissipated in its clocking system. This is largely due to the fact that we did not implement clock gating in this first prototype. All circuits within each processor are clocked continuously—except during idle periods when the oscillator is halted.

The area used by AsAP, shown in Table VI, is the combined area required for all processors including those used for communication. Data for the FIR, 8×8 DCT, and FFT are deduced from measured results of larger applications. We estimated the performance of the JPEG encoder on the TI C62x by using the relative performance of the C62x compared to MIPS processors [28], and a reported similar ARM processor [33].

Figure 15 compares the relative performance and power of an AsAP processor to other processors in Table VI. These comparisons use 8, 8, 13, 9, and 22 AsAP processors—which clearly do not make full use of the chip's 36 processors. Utilizing a larger numbers of processors (through further parallelization) would increase performance further. Nevertheless,

- AsAP achieves 27–275 times higher performance and 96–215 higher energy efficiency than RISC processors (single issue MIPS and ARM);
- compared to a high-end programmable DSP (TI C62x), AsAP achieves 0.8–9.6 times higher performance and 10–75 times higher energy efficiency; and
- compared to ASIC implementations, AsAP achieves performance within a factor of 2–5 and energy efficiency within a factor of 3–50 with an area within a factor of 2.5–3.

Another source of AsAP's high energy efficiency comes from its haltable clock, which is greatly aided by the GALS clocking style. Halting clocks while processors are even momentarily inactive results in power reductions of 53% for the JPEG core and 65% for the 802.11a/g baseband transmitter.

Supply voltage scaling can be used to further improve power savings. Processors dissipate an average of 2.4 mW at a clock rate of 116 MHz using a supply voltage of 0.9 V while executing the described applications.

V. RELATED WORK

There have been many other styles of parallel processors. The key features of the AsAP processor are a small memory, a simple processor, GALS clocking style, and reconfigurable nearest-neighbor mesh network. These features distinguish it from other previous and current parallel processors.

The transputer [44] is a popular parallel processor originally developed in the 1980's. It shares the philosophy of using multiple relatively simple processors to achieve high performance. The transputer is designed for a multiple processor board, where each transputer processor is a complete standalone system. It uses a bit serial channel for inter-processor communication which can support communication of different word lengths to save hardware, but with dramatically reduced communication speeds.

Systolic processors and wavefront processors are two more classic parallel architectures. Systolic processors [45] contain synchronously-operating processors which send and receive data in a highly regular manner [46]. Wavefront array processors [47] are similar to systolic processors but rely on

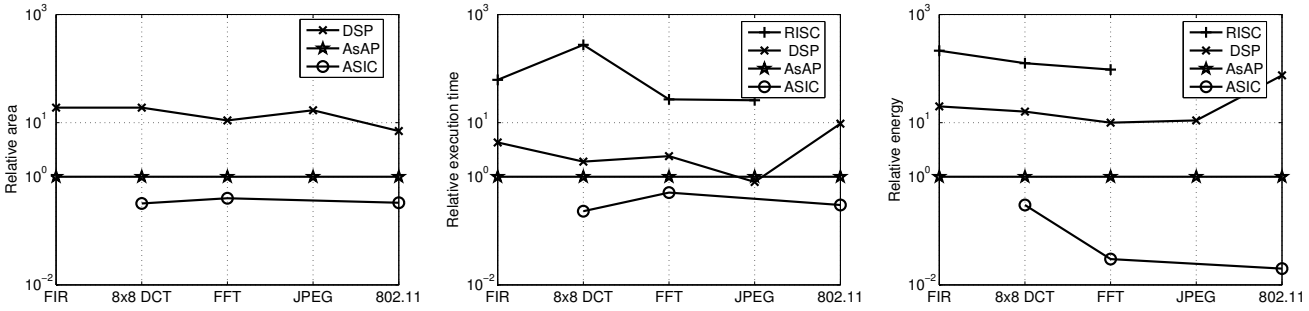

Fig. 15. Relative area, execution time and energy for various implementations of several key DSP kernels and applications. Source data are available in Table VI.

TABLE VII

COMPARISON OF KEY FEATURES OF SELECTED PARALLEL PROCESSOR ARCHITECTURES

Processor	Single Element	Clock Style	Inter-proc Network	Applications
Pleiades [36]	heterogeneous (ASIC + proc.)	handshake GALS	hierarchical network	wireless apps
Picochip [37]	heterogeneous (DSP + coproc.)	synchronous	hierarchical network	wireless apps
Cradle [38]	heterogeneous (RISC + DSP)	synchronous	global bus	multimedia apps
Imagine [39]	ALU cluster	synchronous	hierarchical switch	stream apps
RaPiD [40]	multiple execution units	synchronous	linear array	DSP apps
PipeRench [29]	execution unit stripe	synchronous	linear array	DSP apps
TRIPS [11]	wide-issue processor	synchronous	2-D mesh; dynamic route	all apps
Intel 80-core [26]	VLIW processor	mesochronous	2-D mesh; dynamic route	all apps
Synchroscalar [41]	SIMD processor	rationally related	global interconnect	DSP apps
CELL [15]	SIMD processor	synchronous	high-bandwidth ring bus	multimedia apps
ClearSpeed [42]	SIMD processor	synchronous	N/A	high perf. apps
Sandbridge [43]	SIMD processor with cache	synchronous	N/A	wireless apps
Smart Memories [19]	single-issue proc. with 128 KB mem	synchronous	packet dynamic route	all apps
RAW [10]	single-issue proc. with 128 KB mem	synchronous	static+dynamic route	all apps
AsAP	single-issue proc. with 512 B mem	GALS	2D reconfig. mesh	DSP apps

dataflow properties for inter-processor data synchronization. Previous designs were optimized for simple and single algorithm workloads such as matrix operations [9] and image processing kernels [48].

More parallel processor projects have appeared recently. Table VII compares the key features of other projects to AsAP. Most parallel processors can be easily differentiated by their processing element architectures which can be categorized into three broad types—heterogeneous, multiple execution units (often similar to classic SIMD), and multiple processors (MIMD). A heterogeneous style such as the one used by Pleiades [36], Picochip [37] and Cradle [38] makes the system efficient for specific applications, but results in a non-regular layout and difficult scaling. Recent processors such as RAW [10], CELL [15], TRIPS [11], and Synchroscalar [41], also use MIMD architectures, but can be easily distinguished from AsAP by their larger processing element granularity alone. One of the main reasons for their increased processor granularity is because they target a wider range of applications. Most other projects use a fully synchronous clocking style. Pleiades and FAUST [49] use GALS but with handshaking flow control, which is quite different from the source-synchronous interprocessor communication used in AsAP that is able to sustain full-rate communication of one word per cycle at high clock rates. The Intel 80-core chip [26] employs mesochronous clocking where each processor has the same clock frequency while the clock phase is allowed to vary.

VI. CONCLUSION

The AsAP platform is well-suited for the computation of complex DSP workloads comprised of many DSP tasks, as well as single highly-parallel computationally demanding tasks. By its very nature of having independent clock domains, very small processing elements, and short interconnects, it is highly energy-efficient and capable of high throughput.

Measured results show that on average, AsAP can achieve several times higher performance and 10 times higher energy efficiency than a high performance DSP processor, while utilizing an area more than 10 times smaller.

Areas of interesting future work include: mapping a broader range of applications to AsAP; developing algorithms and hardware for intelligent clock and voltage scaling; automatic software mapping tools to optimize utilization, throughput, and power; C compiler enhancements; connecting large memories when more memory is needed; and automatic fault detection and recovery.

ACKNOWLEDGMENTS

The authors gratefully acknowledge support from Intel, UC Micro, NSF Grant 0430090 and CAREER Award 0546907, SRC GRC Grant 1598, Intellasys, S Machines, MOSIS, Artisan, and a UCD Faculty Research Grant; and thank D. Truong, M. Singh, R. Krishnamurthy, M. Anders, S. Mathew, S. Muroor, W. Li, and C. Chen.

REFERENCES

[1] M. Horowitz and W. Dally, "How scaling will change processor architecture," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2004, pp. 132–133.

[2] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *Proceedings of the IEEE*, pp. 490–504, Apr. 2001.

[3] D. M. Chapiro, *Globally-Asynchronous Locally-Synchronous Systems*, Ph.D. thesis, Stanford University, Stanford, CA, Oct. 1984.

[4] S. Agarwala, T. Anderson, A. Hill, et al., "A 600-MHz VLIW DSP," *IEEE Journal of Solid-State Circuits (JSSC)*, pp. 1532–1544, Nov. 2002.

[5] J. Stinson and S. Rusu, "A 1.5 GHz third generation Itanium 2 processor," in *Design Automation Conference (DAC)*, June 2003, pp. 706–710.

[6] Z. Yu, M. Meeuwsen, R. Apperson, et al., "An asynchronous array of simple processors for DSP applications," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2006, pp. 428–429.

[7] N. Bindal et al., "Scalable sub-10ps skew global clock distribution for a 90nm multi-GHz IA microprocessor," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2003, pp. 346–347.

[8] R. Apperson, Z. Yu, M. Meeuwsen, T. Mohsenin, and B. Baas, "A scalable dual-clock FIFO for data transfers between arbitrary and haltable clock domains," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 15, no. 10, pp. 1125–1134, Oct. 2007.

[9] S. Y. Kung, "VLSI array processors," in *IEEE ASSP Magazine*, July 1985, pp. 4–22.

[10] M. Taylor et al., "A 16-issue multiple-program-counter microprocessor with point-to-point scalar operand network," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2003, pp. 170–171.

[11] S. Keckler et al., "A wire-delay scalable microprocessor architecture for high performance systems," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2003, pp. 168–169.

[12] J. Glossner, J. Moreno, M. Moudgill, et al., "Trends in compilable DSP architecture," in *IEEE Workshop on Signal Processing Systems (SiPS)*, Oct. 2000, pp. 181–199.

[13] IEEE Computer Society, "Wireless LAN medium access control (MAC) and physical layer (PHY) specifications: High speed physical layer in the 5 GHz band," in *Standard for Information Technology*. Institute of Electrical and Electronics Engineers, 1999.

[14] M. J. Meeuwsen, O. Sattari, and B. M. Baas, "A full-rate software implementation of an IEEE 802.11a compliant digital baseband transmitter," in *IEEE Workshop on Signal Processing Systems (SiPS)*, Oct. 2004, pp. 124–129.

[15] D. Pham et al., "The design and implementation of a first-generation CELL processor," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2005, pp. 184–185.

[16] S. Naffziger et al., "The implementation of a 2-core multi-threaded Itanium family processor," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2005, pp. 182–183.

[17] J. Oliver et al., "Tile size selection for low-power tile-based architecture," in *ACM Computing Frontiers*, May 2006, pp. 83–94.

[18] B. Baas, "A parallel programmable energy-efficient architecture for computationally-intensive DSP systems," in *Asilomar Conference on Signals, Systems and Computers*, Nov. 2003, pp. 2185–2189.

[19] K. Mai et al., "Smart memories: A modular reconfigurable architecture," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2000, pp. 161–171.

[20] J. Sungtae et al., "Energy characterization of a tiled architecture processor with on-chip network," in *International Symposium on Low Power Electronics and Design (ISLPED)*, Aug. 2003, pp. 424–427.

[21] A. A. Bright et al., "Creating the BlueGene/L supercomputer from low-power SOC ASICs," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2005, pp. 188–189.

[22] A. S. Leon et al., "A power-efficienct high-throughput 32-thread SPARC processor," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2006, pp. 98–99.

[23] Texas Instruments, "DSP platforms benchmarks," Tech. Rep., `http://www.ti.com/`.

[24] Berkeley Design Technology, *Evaluating DSP Processor Performance*, Berkeley, CA, USA, 2000.

[25] the Embedded Microprocessor Benchmark Consortium, *Data sheets*, www.eembc.org, 2006.

[26] S. Vangal, J. Howard, G. Ruhl, et al., "An 80-tile 1.28TFLOPS network-on-chip in 65nm CMOS," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2007, pp. 98–99.

[27] Z. Yu et al., "Performance and power analysis of globally asynchronous locally synchronous multi-processor systems," in *IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, Mar. 2006, pp. 378–384.

[28] C. Kozyrakis et al., "Vector vs. superscalar and vliw architectures for embedded multimedia benchmarks," in *Micro*, Nov. 2002, pp. 283–289.

[29] H. Schmit et al., "PipeRench: A virtualized programmable datapath in 0.18 micron technology," in *IEEE Custom Integrated Circuits Conference (CICC)*, May 2002, pp. 63–66.

[30] B. Gorjiara et al., "Custom processor design using NISC: a case-study on DCT algorithm," in *ESTIMedia*, Sept. 2005, pp. 55–60.

[31] M. Matsui et al., "A 200 MHz 13 mm^2 2-d DCT macrocell using sense-amplifying pipeline flip-flop scheme," *IEEE Journal of Solid-State Circuits (JSSC)*, pp. 1482–1490, Dec. 1994.

[32] K. Maharatna et al., "A 64-point fourier transform chip for high-speed wireless LAN application using OFDM," *IEEE Journal of Solid-State Circuits (JSSC)*, pp. 484–493, Mar. 2004.

[33] T. Lin and C. Jen, "Cascade – configurable and scalable DSP environment," in *IEEE International Symposium on Circuits and Systems (ISCAS)*, May 2002, pp. 26–29.

[34] M. Tariq et al., "Development of an OFDM based high speed wireless LAN platform using the TI C6x DSP," in *IEEE International Conference on Communications (ICC)*, Apr. 2002, pp. 522–526.

[35] J. Thomson et al., "An Integrated 802.11a Baseband and MAC Processor," in *IEEE International Solid-State Circuits Conference (ISSCC)*, 2002, vol. 45, pp. 126–127, 451.

[36] H. Zhang et al., "A 1-V heterogeneous reconfigurable DSP IC for wireless baseband digital signal processing," *IEEE Journal of Solid-State Circuits (JSSC)*, pp. 1697–1704, Nov. 2000.

[37] R. Baines et al., "A total cost approach to evaluating different reconfigurable architectures for baseband processing in wireless receivers," *IEEE Communication Magazine*, pp. 105–113, Jan. 2003.

[38] Cradle Technologies, "Multiprocessor DSPs: Next stage in the evolution of media processor DSPs," Tech. Rep., `http://www.cradle.com/`.

[39] B. Khailany et al., "VLSI design and verification of the imagine processor," in *IEEE International Conference on Computer Design (ICCD)*, Sept. 2002, pp. 289–294.

[40] D. C. Cronquist et al., "Architecture design of reconfigurable pipelined datapaths," in *Advanced research in VLSI (ARVLSI)*, Mar. 1999, pp. 23–40.

[41] J. Oliver et al., "Synchroscalar: A multiple clock domain, power-aware, tile-based embedded processor," in *Proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2004, pp. 150–161.

[42] ClearSpeed, "CSX600: Advanced product," Tech. Rep., `http://www.clearspeed.com/`.

[43] Sandbridge, "The sandbridge sandblaster convegence platform," Tech. Rep., `http://www.sandbridgetech.com/`.

[44] C. Whitby-Strevens, "Transputers-past, present and future," *IEEE Micro*, pp. 16–19, Dec. 1990.

[45] H. T. Kung, "Why systolic architectures?," in *Computer Magazine*, Jan. 1982, pp. 37–46.

[46] H. T. Kung, "Systolic communication," in *International Conference on Systolic Arrays*, May 1988, pp. 695–703.

[47] S. Kung et al., "Wavefront array processor: Language, architecture, and applications," *IEEE Transactions on Computers*, vol. C-31, no. 11, Nov. 1982.

[48] U. Schmidt and S. Mehrgardt, "Wavefront array processor for video applications," in *IEEE International Conference on Computer Design (ICCD)*, Sept. 1990, pp. 307–310.

[49] D. Lattard, E. Beigne, C. Bernard, et al., "A telecom baseband circuit based on an asynchronous network-on-chip," in *IEEE International Solid-State Circuits Conference (ISSCC)*, Feb. 2007, pp. 258–259.

Zhiyi Yu received the B.S. and M.S. degrees in electrical engineering from Fudan University, Shanghai, China, in 2000 and 2003, respectively, and the Ph.D degree in electrical and computer engineering from University of California, Davis, in 2007.

Dr. Yu is currently a Hardware Engineer with IntellaSys Corporation, headquartered in Cupertino, CA. His research interests include high-performance and energy-efficient digital VLSI design, architectures, and processor interconnects, with an emphasis on many-core processors. He was a key designer of the 36-core Asynchronous Array of simple Processors (AsAP) chip, and one of the designers of the 150+ core second generation computational array chip.

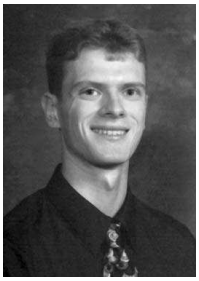

Michael J. Meeuwsen received the B.S. degrees with honors in electrical engineering and computer engineering (both summa cum laude) from Oregon State University, Corvallis, and the M.S. in electrical and computer engineering from the University of California, Davis.

He is currently a Hardware Engineer with Intel Digital Enterprise Group, Hillsboro, OR, where he works on CPU hardware design. His research interests include digital circuit design and IEEE 802.11a/g algorithm mapping.

Ryan W. Apperson received the B.S. in electrical engineering (Magna Cum Laude) from the University of Washington, Seattle, and the M.S. degree in electrical and computer engineering from the University of California, Davis.

He is currently an IC Design Engineer with Boston Scientific CRM Division, Redmond, WA. His research interests include multiclock domain systems and SRAM design.

Omar Sattari received the B.S. and M.S. degrees in electrical and computer engineering from the University of California, Davis.

He is currently a Software Engineer at CornerTurn. His research interests include FFT and DSP algorithms and digital hardware design.

Michael Lai received the B.S. and M.S. degrees in electrical and computer engineering from the University of California, Davis.

He is currently a Design Engineer at Altera Corporation working on next generation transceiver products. His research interests include the design of high-speed arithmetic units and control.

Jeremy W. Webb received the B.S. degree in electrical and computer engineering from the University of California, Davis.

He is currently a M.S. student in electrical and computer engineering at the University of California, Davis, and a hardware engineer at Centellax. His research interests include high-speed board design and system interfacing.

Eric W. Work received the B.S. degree from the University of Washington, and the M.S. degree in electrical and computer engineering from the University of California, Davis.

He is currently a Software Engineer at S Machine Corporation. His research interests include the mapping of arbitrary task graphs to processor networks and software tool flow.

Tinoosh Mohsenin received the B.S. degree in electrical engineering from Sharif University, Tehran, Iran, and the M.S. degree in electrical and computer engineering from Rice University, Houston, TX. She is currently pursuing the Ph.D. degree in electrical and computer engineering from the University of California, Davis.

She is the designer of the Split-Row and Multi-Split-Row Low Density Parity Check (LDPC) decoding algorithms. Her research interests include energy efficient and high performance signal processing and error correction architectures including multi-gigabit full-parallel LDPC decoders and many-core processor architecture design.

Bevan M. Baas received the B.S. degree in electronic engineering from California Polytechnic State University, San Luis Obispo, in 1987, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1990 and 1999, respectively.

In 2003 he became an Assistant Professor with the Department of Electrical and Computer Engineering at the University of California, Davis. He leads projects in architecture, hardware, software tools, and applications for VLSI computation with an emphasis on DSP workloads. Recent projects include the Asynchronous Array of simple Processors (AsAP) chip, applications, and tools; low density parity check (LDPC) decoders; FFT processors; viterbi decoders; and H.264 video codecs.

From 1987 to 1989, he was with Hewlett-Packard, Cupertino, CA, where he participated in the development of the processor for a high-end minicomputer. In 1999, he joined Atheros Communications, Santa Clara, CA, as an early employee and served as a core member of the team which developed the first IEEE 802.11a (54 Mbps, 5 GHz) Wi-Fi wireless LAN solution. During the summer of 2006 he was a Visiting Professor in Intel's Circuit Research Lab.

Dr. Baas was a National Science Foundation Fellow from 1990 to 1993 and a NASA Graduate Student Researcher Fellow from 1993 to 1996. He was a recipient of the National Science Foundation CAREER Award in 2006 and the Most Promising Engineer/Scientist Award by AISES in 2006. He is an Associate Editor for the IEEE JOURNAL OF SOLID-STATE CIRCUITS and has served as a member of the Technical Program Committee of the IEEE International Conference on Computer Design (ICCD) in 2004, 2005, and 2007. He also serves as a member of the Technical Advisory Board of an early stage technology company.