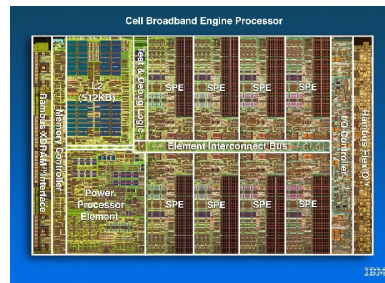


A Survey of Multi-Core Coarse-Grained Reconfigurable Arrays for Embedded Applications



Justin L. Tripp, Jan Frigo, Paul Graham
Los Alamos National Laboratory
Los Alamos, NM 87545

September 2007

Embedded Applications

- Traditionally handled by DSPs and FPGAs
- DSPs
 - ★ high clocks speeds, low power
 - ★ programmability
 - ★ low flexibility
- FPGAs
 - ★ multilevel parallelism (bit, operation, etc.)
 - ★ high flexibility
 - ★ challenging programmability
- Factors - power, speed, configurability, programmability
- Applications continue to demand more computation.

Coarse-Grained Reconfigurable Array

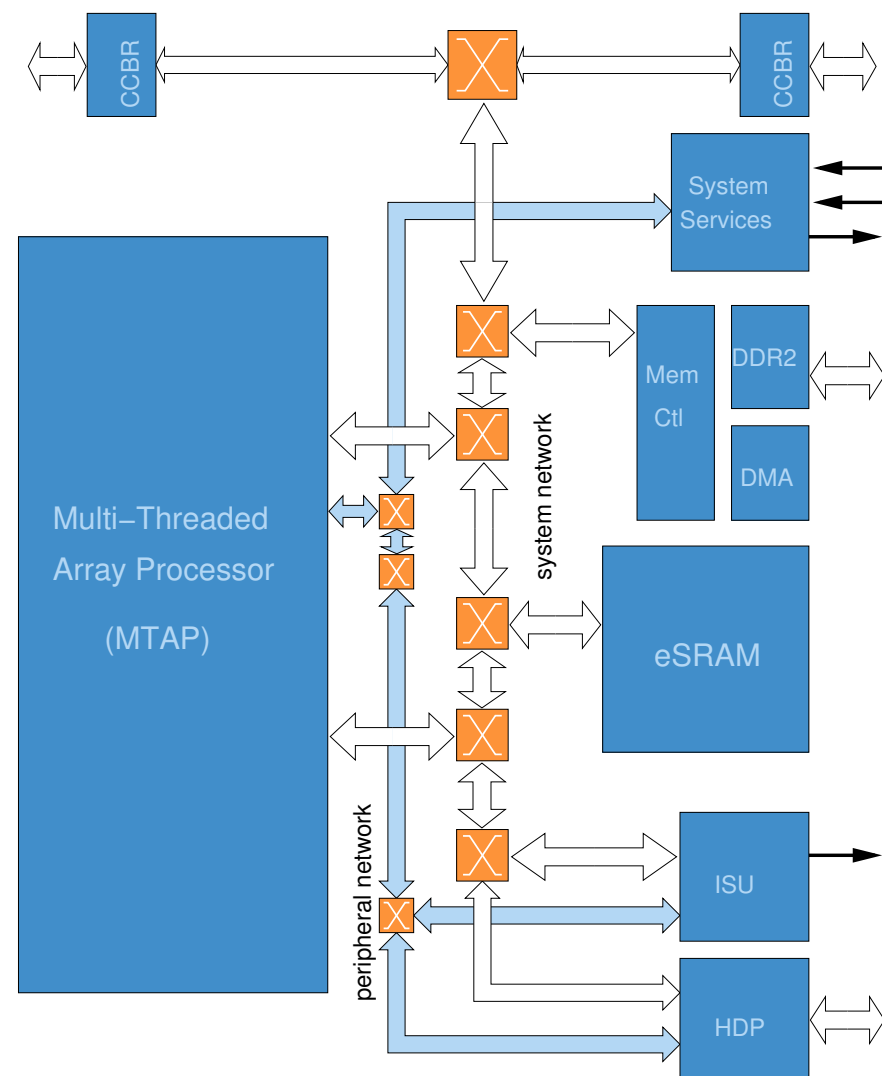
- Expand the solution space for power, speed, configurability, embeddability and programmability
- Higher Performance for lower power
- Processor Centric
 - ★ Individual Processors connected with by a network.
 - ★ Examples: IBM Cell BE, Clearspeed CSX600, GPUs, etc.
- Medium-Grained Arrays
 - ★ Arrays with the basic unit of computation higher level than a gate.
 - ★ Examples: Exilent's reconfigurable arithmetic processors, MathStar's FPOA, etc.

Programmer Productivity

- “[I]ndustry average for code production is 8-20 lines of correct code per day” – (Steve McConnell, Code Complete, 1993)
- If you are writing in a Hardware Description Language or Assembly, you are going to produce fewer lines of debugged code per day.
- High-level languages increase the expressiveness of the code and potentially reduce the amount of code necessary to express a given idea.
- For some applications the highest performance is not as important as providing the product by a certain deadline.

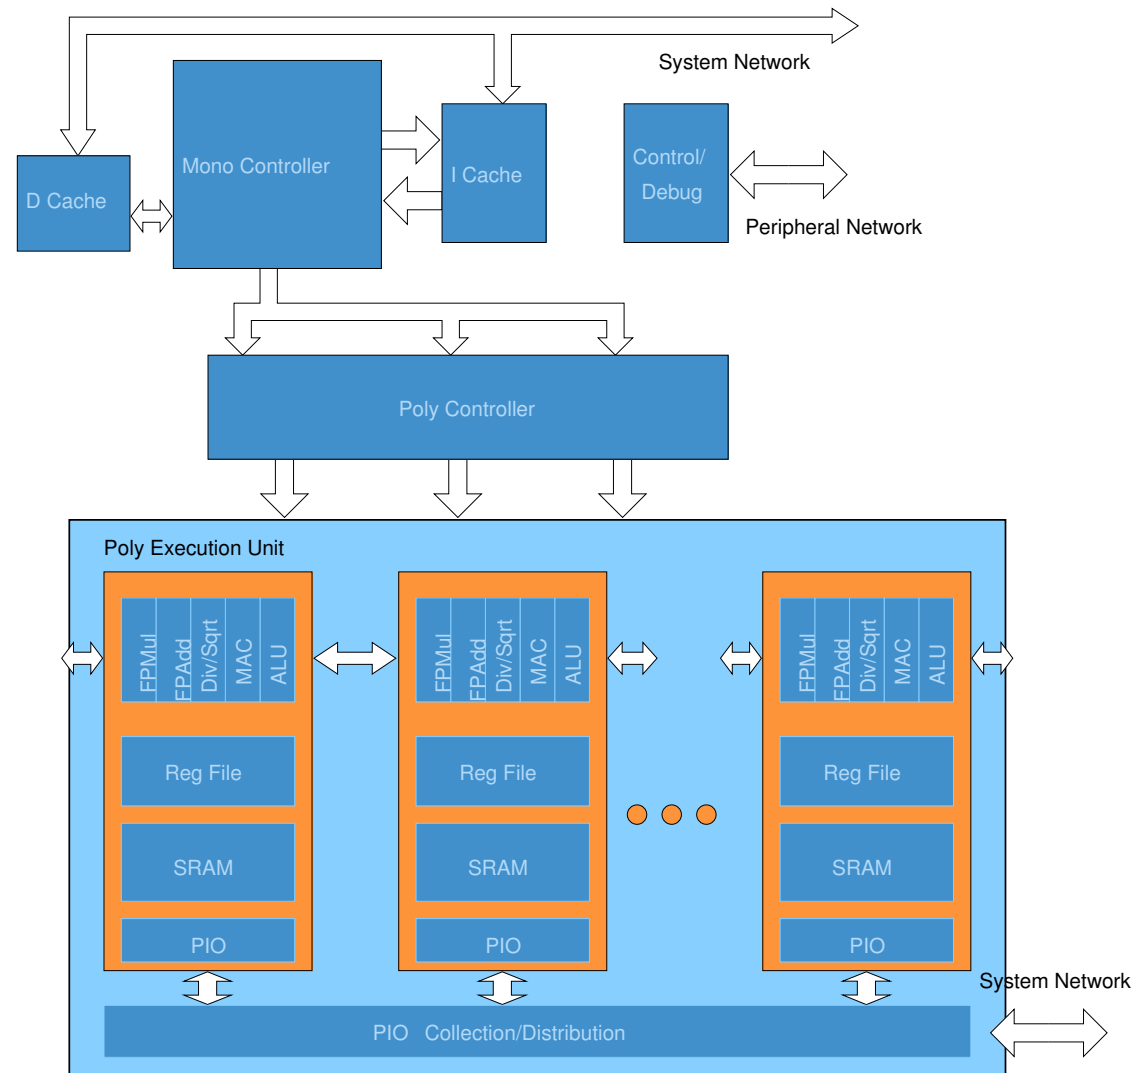
ClearSpeed CSX600

- Floating Point Array (FPA) with a RISC control processor which issues instructions in a SIMD fashion.
- RISC processor and FPA run at 250MHz and there are 96 Processor Elements.
- Low Power - providing 25GFlops at 10Watts.
- CSX600 contains MTAP processor, DRAM interface, high speed I/O, embedded 128kB SRAM.



ClearSpeed MTAP

- 96 Processing Elements
- 5 Processing Unit per PE (FPMul, FPAdd, int Div/Sqrt, 16bit int MAC, int ALU)
- Each PE has 6kB SRAM and 128 byte Register file.
- PEs can transfer 1 16bit word in 2 cycles



Complex FFT Results

Complex FFT of 1024, 2048, and 4096 pts using 2 CSX600 Chips.

Dimension/Size	Direction	Type	FFT/sec	usec/FFT	GFlops/sec
1D 1024	forward	single complex	157242.28	6.360	8.05
	backward	single complex	156870.43	6.375	8.03
	forward	double complex	78168.97	12.79	4.00
	backward	double complex	77460.23	12.91	3.97
1D 2048	forward	single complex	77180.89	12.96	8.69
	backward	single complex	70430.39	14.20	7.93
	forward	double complex	38896.51	25.71	4.38
	backward	double complex	38831.41	25.75	4.37
1D 4096	forward	single complex	31986.69	31.26	7.86
	backward	single complex	29768.52	33.59	7.32
	forward	double complex	19254.78	51.94	4.73
	backward	double complex	19225.21	52.02	4.72

Two CSX600 Chip board requires 25 Watts of power. FPGAs can achieve 13 μ sec for a 4k 16bit FFT at 5 Watts.

Stream Benchmark

The STREAM benchmark is a simple synthetic benchmark program that measures sustainable memory bandwidth and the corresponding computation rate for simple vector kernels.

Name	Kernel
Copy	$a(i) = b(i)$
Scale	$a(i) = q \times b(i)$
Sum	$a(i) = b(i) + c(i)$
Triad	$a(i) = b(i) + q \times c(i)$

STREAM results (MB/s)

Processor	Copy	Scale	Add	Triad
ClearSpeed DRAM to RISC	34.3	32.9	31.5	30.7
ClearSpeed PE SRAM	1327.4	1276.1	1808.8	1743.4
Per PE	13.9	13.2	18.8	18.2
ClearSpeed DRAM to PE	587.7	22.6	27.5	24.2
Per PE	6.1	0.2	0.3	0.3
Opteron	2029.4	2110.0	2486.1	2539.0
60MHz Pentium	37.2	62.1	61.3	58.5

Advection Simulation

Advection describes the transport of a scalar quantity in a vector field such as the movement of silt in a river.

$$F_{new} = \frac{F_{old} \times V_{old} + \sum (F' dv)}{v_{new}}$$

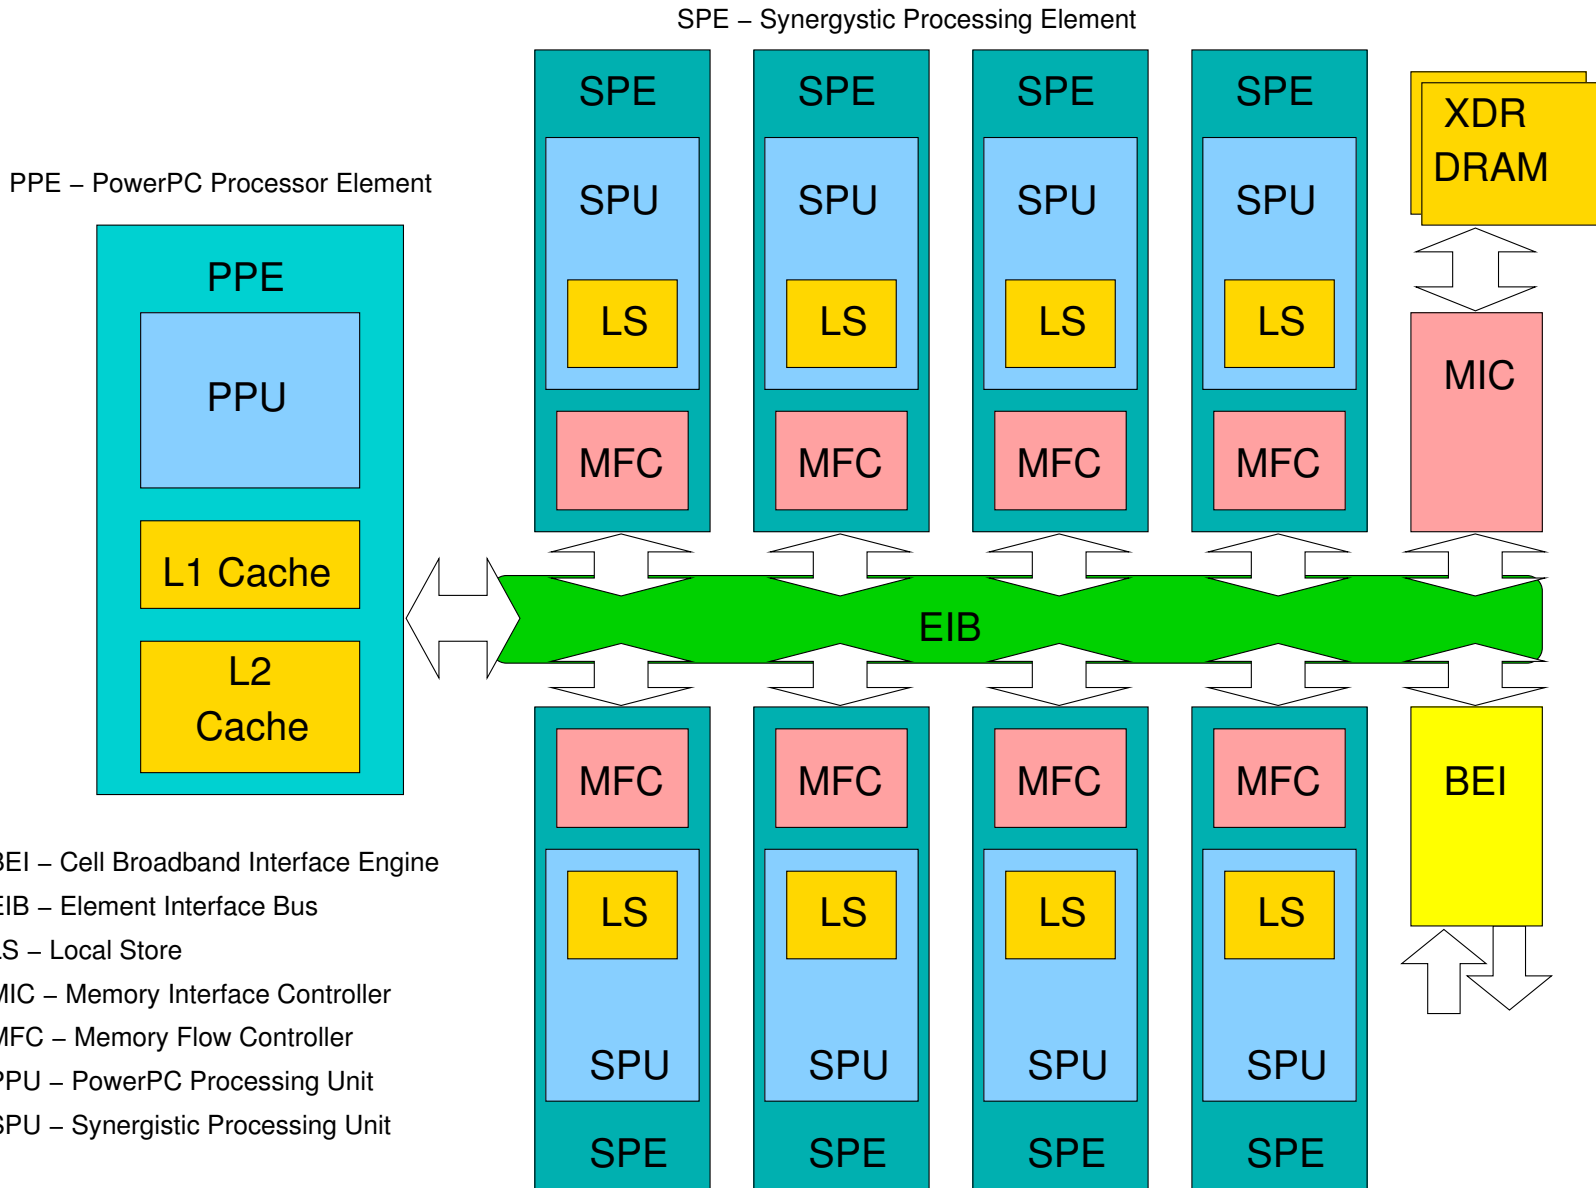
Implementation	Execution Time
Simple Memory Copies	1m 55s
Asynchronous Memory Copies	1m 37s
Triple Buffered	1m 41s
Swazzles Only	0m 46s
Opteron	0m 43s

The difference in implementations lies in how inter-PE communication is handled. In the first set of examples, communication is completed between PE and shared DRAM. In the final implementation, all communication is locally between PEs as PE-to-PE swazzles.

IBM Cell Broadband Engine

- Designed to overcome problems with current processor design: memory latency and bandwidth, power, and diminishing returns achieved by increasing pipelining depth.
- Single-chip multiprocessor with 9 processors share one coherent memory
- Programmable local caches
- Fast communication between processors (205GB/s ring bandwidth, 25GB/s XDR memory)
- Low-power CMOS SOI technology
- Short pipelines, SIMD instructions, no branch prediction

Cell Architecture



Complex FFT Results


Results for complex single precision (SP) floating point. Throughput reported in μ secs.

	256	1024	2048	4096	8192	64K
Cell (3.2GHz)	0.53	2.4	5.07	11.37	23.83	57.74
Xeon64 (3.6GHz)	1.18	6.4	15.7	37.6	79.3	864
Opteron (2.4GHz)	2.15	10.62	23.37	54.01	130.4	1725

This compares favorably to the FPGA result of 13.04 μ secs for a 4K 16 bit FFT. But the power for the Cell processing board is 225W.

Matched Filter

The matched filter is an important kernel in the processing of hyperspectral data. The filter enables researchers to sift useful data from instruments that span large frequency bands and can produce Gigabytes of data in seconds.



Datacube and layers of 2D image for each spectral band of interest.

Implementation	Speedup over CPU	Time per Signature
FPGA (v2pro40)	3.91	0.78 sec
GPU (nv 7950)	3.1	1.0 sec
Cell (3GHz)	8.0	0.38 sec
CPU (3GHz Xeon)	1	3.05 sec

Stretch S5000

- Overview

- ★ Tensilica Xtensa RISC processor Core
- ★ Stretch Inst. Set Extension Fabric
- ★ 300 MHz Clock Rate

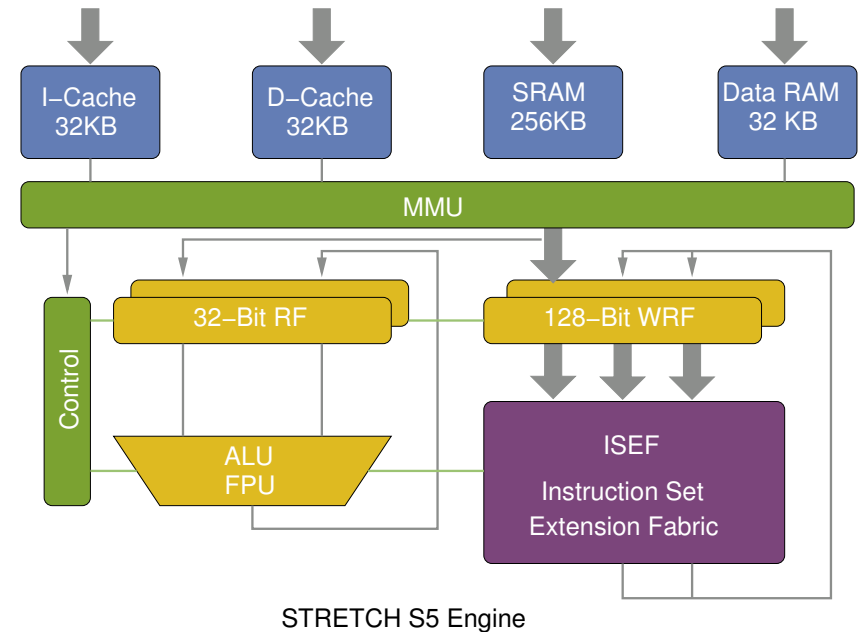
- Internal Cache - I, D, SRAM, Data RAM

- On-chip MMU, 24 Ch. of DMA, FPU

- Wide Register File (32 by 128 bit)

- ISEF


- ★ 3 inputs/2 outputs
- ★ Pipelined, interlocked



- ★ 32 16-bit MACs and 256 ALUs
- ★ Bit-sliced for arbitrary width
- ★ Allow control logic
- ★ Allow processor internal state

ISEF Architecture

- Array of 64-bit ALUs
 - ★ May be broken at 4-bit and 1-bit boundaries
 - ★ Conditional ALU operation: $Y = C ? (A \text{ op}_1 B) : (A \text{ op}_2 B)$
 - ★ Registers to implement state, pipelining
- Array of 4×8 multipliers, cascadable to 32×32
 - ★ Programmable pipelining
- Programmable Routing Fabric
- Execution clock divided from processor clock: 1:1, 1:2, 1:3, 1:4
- Up to 16 instructions per ISEF configuration
- Many configurations per executable
- Reconfiguration
 - ★ User directed or on demand
 - ★ $100 \mu\text{sec}$ for complete reconfiguration



Performance Stretch S5000 Processor

	Data	Throughput (μs)	Freq. (MHz)	Power (W)
1k FFT (radix 2)	16-bit	12.8	300	3.5
4k FFT (radix 2)	16-bit	61.4	300	3.5
8k complex FFT	16-bit	260.6	300	3.5
64 Tap FIR	16-bit	2.55	300	3.5

Image Processing Application

- Application

- ★ Real-time anomaly detection for a process that machines metal parts
- ★ 4000 frames/sec, 100 to 200 frames for cross-correlation
- ★ 100×110 image size, 8-bit unsigned data

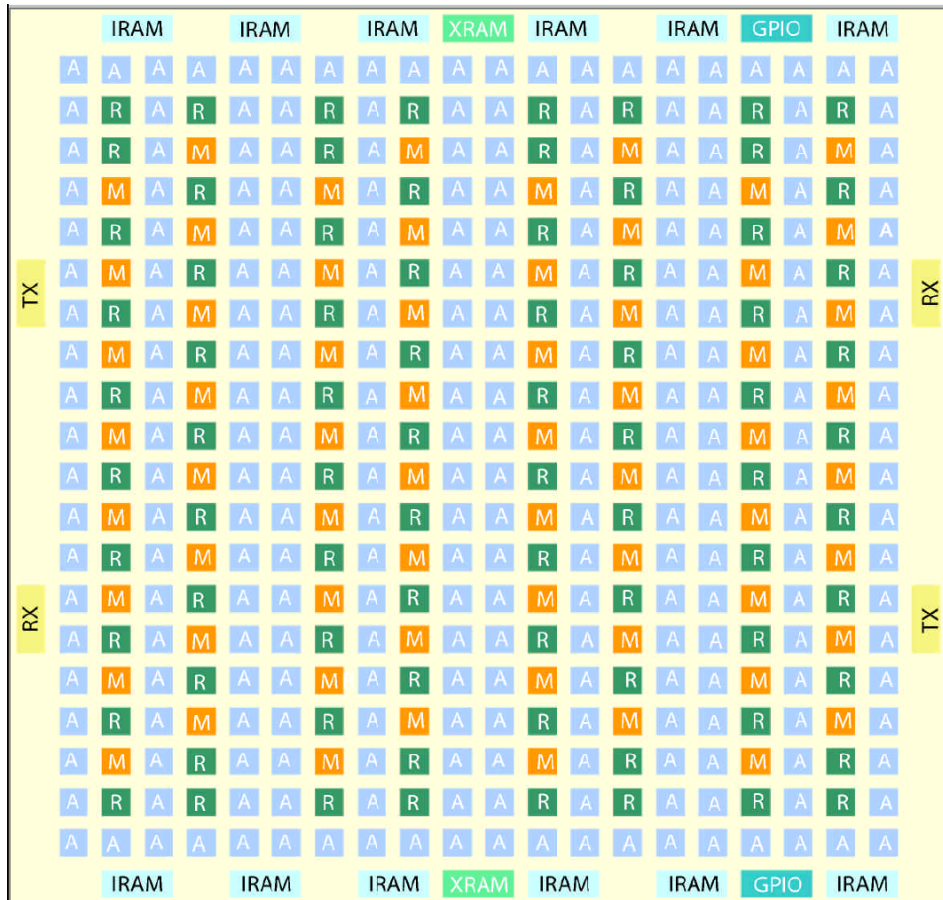
- Algorithm

- ★ High Pass Filter - to enhance edges
- ★ Interframe differencing - isolate moving particles
- ★ Cross-correlation - track movement of particles
- ★ Norm of the convolution

	Stretch S5 (300 MHz)	Intel Xeon (3.06 GHz)	Relative Performance
Cycles per Frame (10^6)	1.86	17.75	9.54×
Throughput (fps)	161	172	0.936×
Power (W)	3.2	85	26.6×

MathStar FPOA


- 400 Silicon Objects, 1 GHz max clock



- ★ 256 16-bit ALUs
- ★ 64 16×16 MACs
- ★ 80 Register Files (RF) - FIFO or configurable RAM
- ★ 12 2K×76 bit memory blocks (IRAM)

- Communication Network

- ★ Nearest Neighbor (0 latency)
- ★ Party Line - 3 objects away (1 clock cycle)

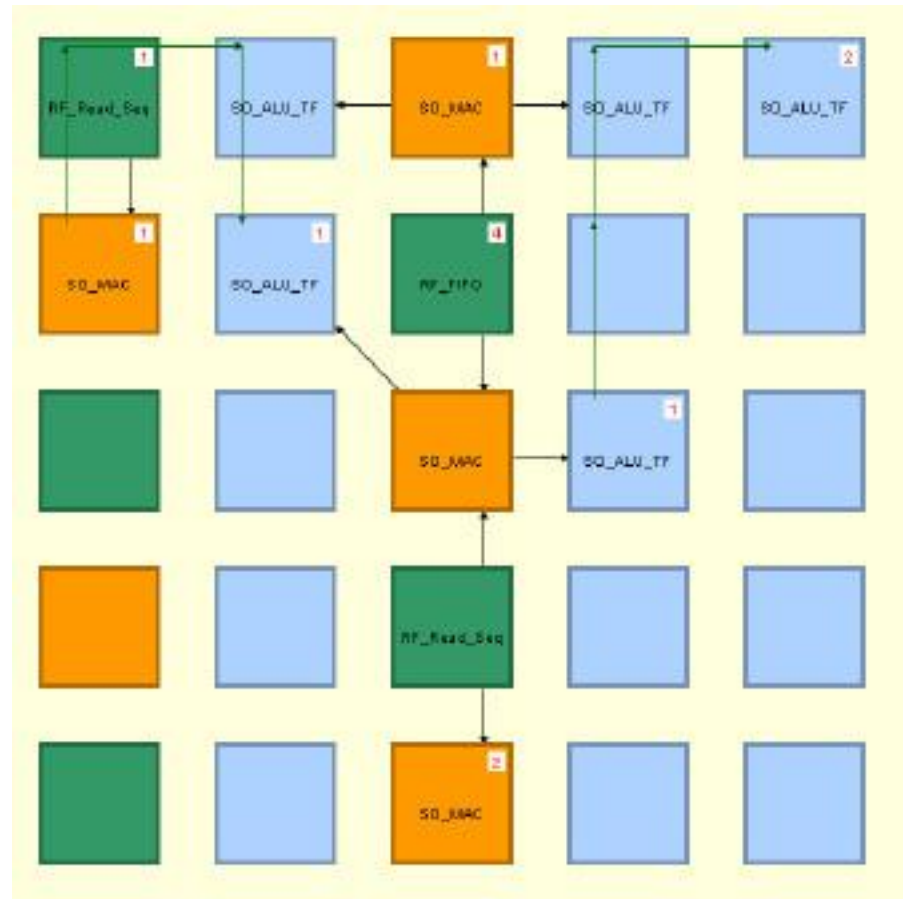


Silicon Object Characteristics

- ALU
 - ★ 16 bit data bus, 32 operations
 - ★ 5 instruction wide VLIW
 - ★ 5 bits control, control logic
- RF
 - ★ 128 Byte, 16 bits wide
 - ★ Dual Port RAM or FIFO
- MAC
 - ★ 16×16 Multiply
 - ★ 40bit accumulate
- IRAM - $2K \times 76$ bits

Eight Tap FIR

- FPOA Resources
 - ★ 5 ALUs
 - ★ 4 MACs
 - ★ 3 Register Files
- Performance
 - ★ 1 GHz
 - ★ 6.7 - 7.3 Watts
 - ★ Throughput - 1000 M samples/s - latency 8 cycles



4×4 Sum of Absolute Difference Operator

- Resources

- ★ 21 ALUs

- ★ 1 Register File

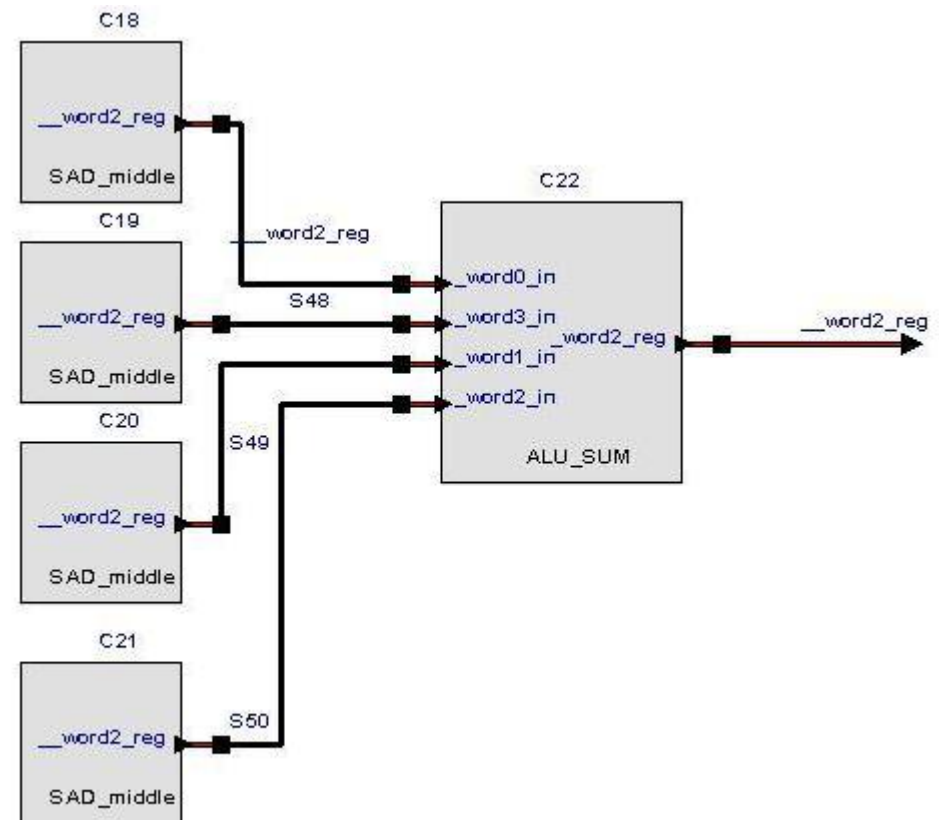
- ★ 1 IRAM

- Performance

- ★ 1 GHz

- ★ 8 Watts


- ★ Throughput - 8 cycles



Comparision



CSX600



Cell BE



Stretch S5



Math* FPOA

Advantages

- 32/64 FP
- Power (10W)
- C Development
- SIMD exec.

- 32/64 int, 32 FP
- 3GHz clock
- Program. Cache
- SIMD exec.

- 32 int/fp
- C/C++ Devel.
- Speed Up/Watt
- Program. Cache

- 16 int, no FP
- 1.78 Mbit RAM
- 1GHz clock

Disadvantages

- Immature Tools
- 250MHz
- 96 way SIMD
- Low Comm. BW

- Program. Cache
- 180W Power
- 256k MEM
- Coarse Grained

- No ISEF RAM
- 300Mhz
- Program. Cache

- Immature Tools
- Manual mapping
- 10-60 Watts

Conclusions

- Nothing is free: Power, Performance, Development Time, Embeddability.
- Each of the architectures plot a different point in a multidimensional space.
- FPGAs, MathStar are high performance and low power, but cost in development time.
- Cell is appropriate when float-point fidelity required and performance is a premium.
- Clearspeed supports floating-point and development time is low.
- Stretch work well on integer applications and where power is a premium.