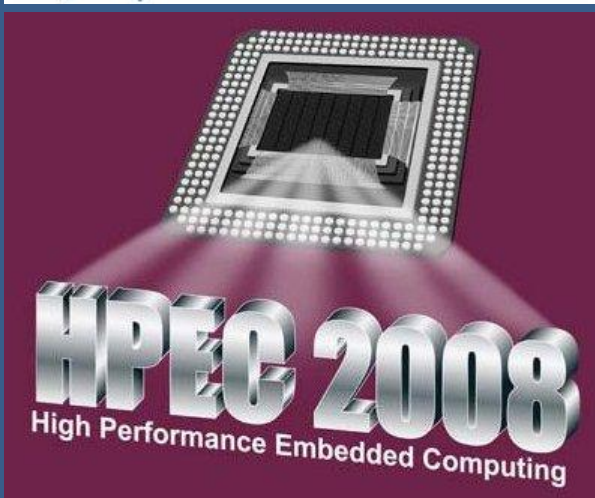




# Fixed and Reconfigurable Multi-Core Device Characterization for HPEC



Jason Williams  
Alan D. George  
Justin Richardson  
Kunal Gosrani  
Siddarth Suresh

**NSF CHREC Center  
ECE Department, University of Florida**

# Outline

- Background
- RC Taxonomy
- Reconfigurability Factors
- Computational Density Metrics
- Internal Memory Bandwidth Metric
- Results & Analysis
- Future Work
- Conclusions



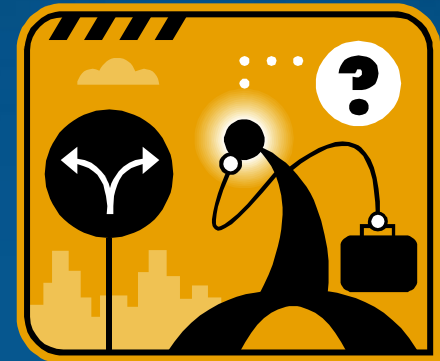
# Background

- Moore's law continues to hold true, transistor counts doubling every 18 months
  - But can no longer rely upon increasing clock rates ( $f_{clk}$ ) and instruction-level parallelism (ILP) to meet computing performance demands
- How to best exploit ever-increasing on-chip transistor counts?
  - Architecture Reformation: Multi- & many-core (MC) devices are new technology wave
  - Application Reformation: focus on exploiting explicit parallelism in these new devices

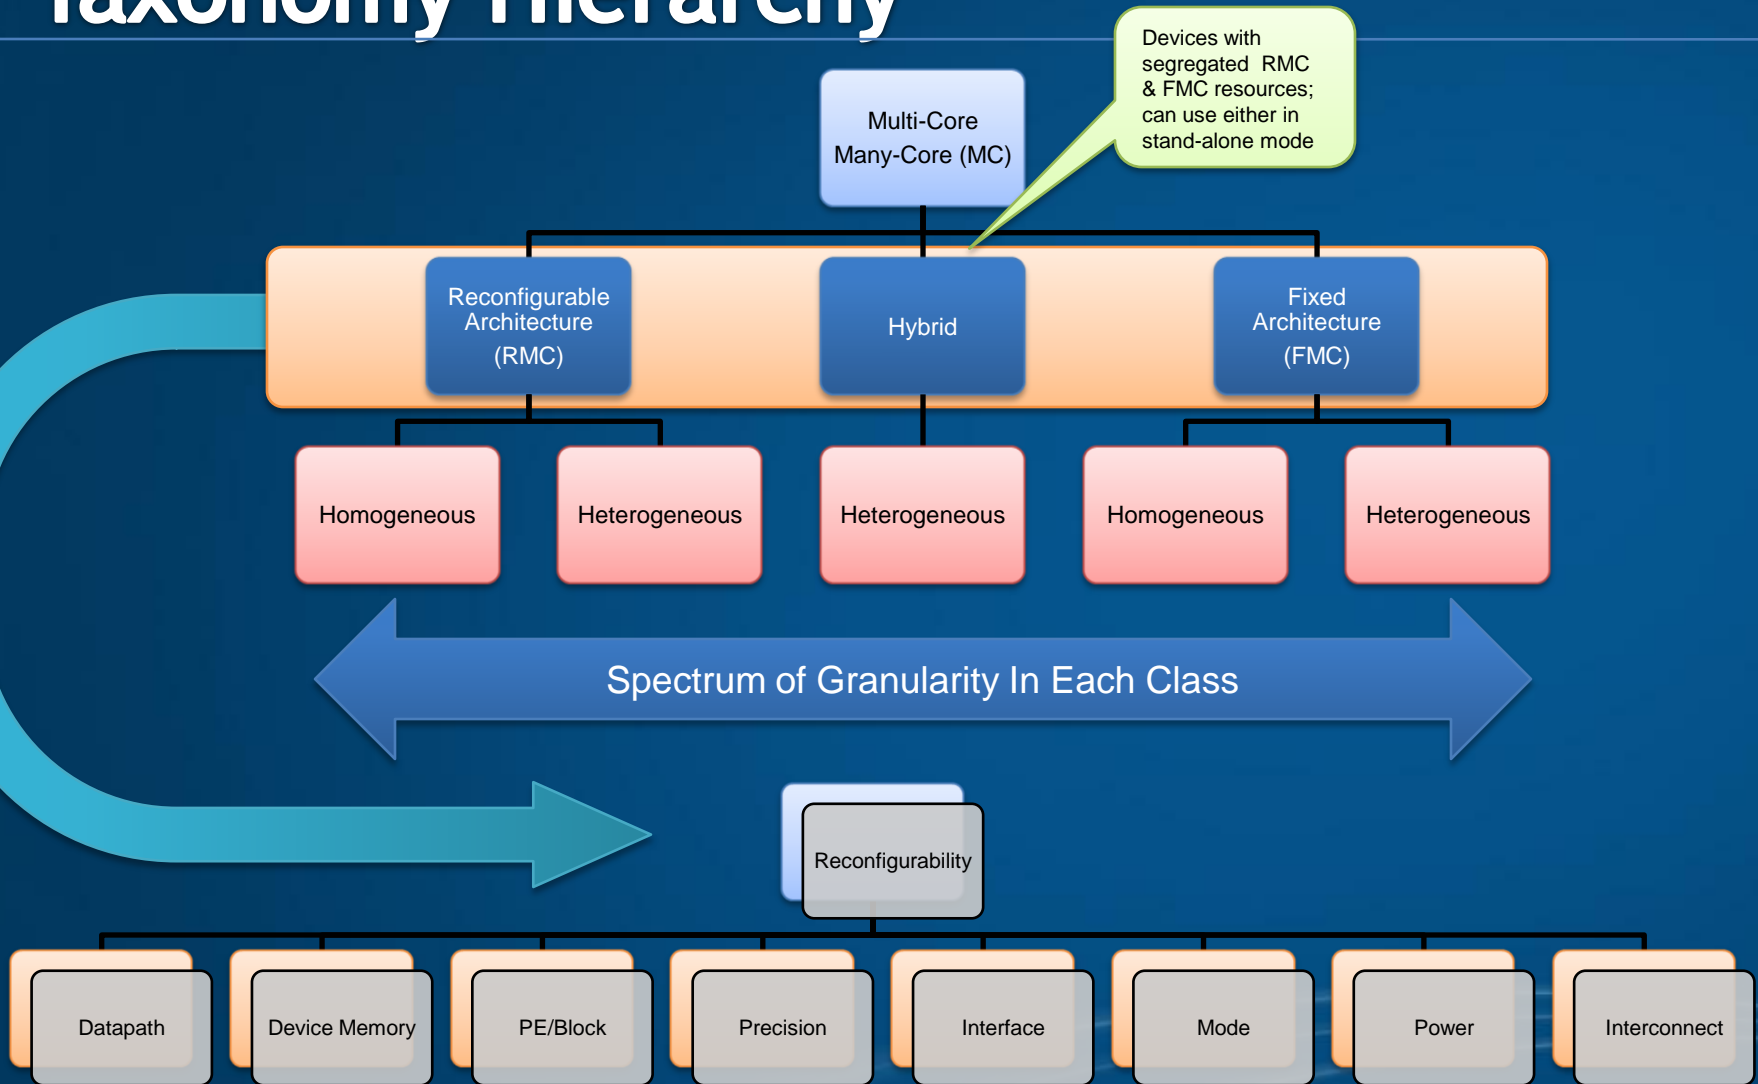


# Background

- What MC architecture options are available?
  - Fixed MC: fixed hardware structure, cannot be changed post-fab
  - Reconfigurable MC: can be adapted post-fab to changing problem req's
- How to compare disparate device technologies?
  - Need for taxonomy & device analysis early in development cycle
  - Challenging due to vast design space of FMC and RMC devices
  - We are developing a suite of metrics; two are focus of this study:
    - Computational Density per Watt captures computational performance and power consumption, more relevant for HPEC than pure performance metrics
    - Internal Memory Bandwidth describes device's on-chip memory access capabilities

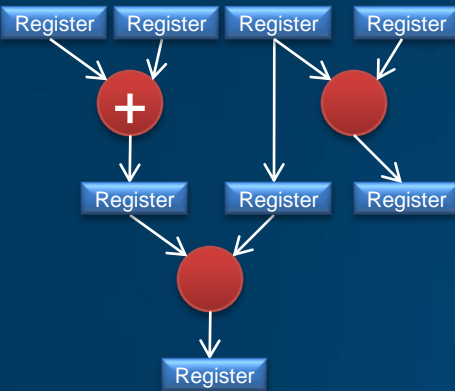


# Taxonomy Hierarchy

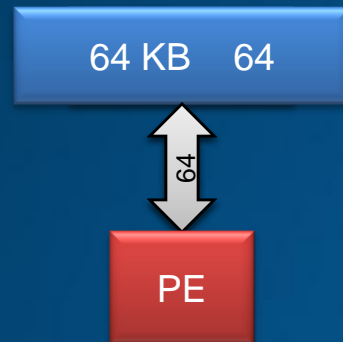


# Reconfigurability Factors

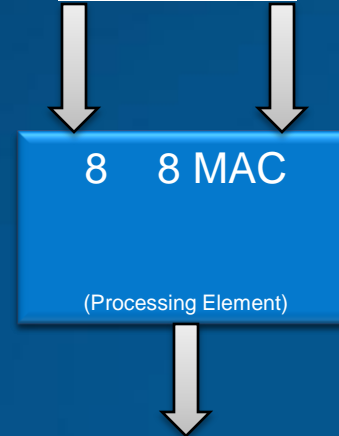
## Datapath



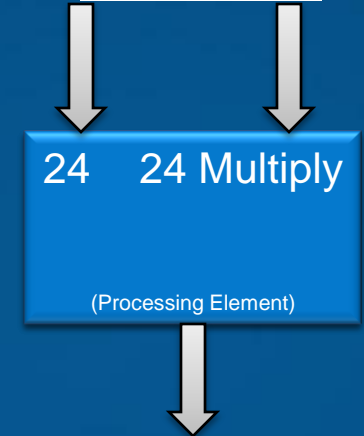
## Device Memory



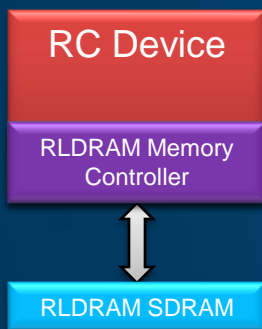
## PE/Block



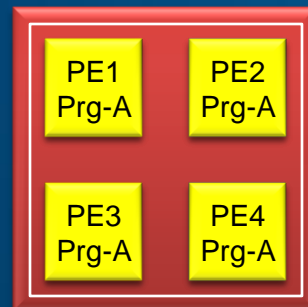
## Precision



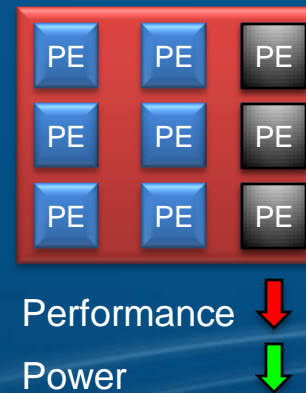
## Interface



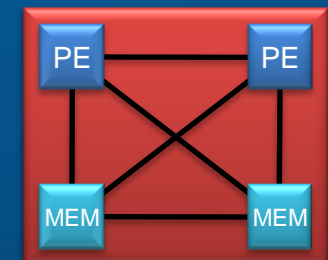
## Mode



## Power



## Interconnect



# Metric Overview

- Metric Description

- Computational Density (CD)
  - Measure of computational performance across range of parallelism, grouped by process technology
- Computational Density per Watt (CDW)
  - CD normalized by power consumption
- Internal Memory Bandwidth (IMB)
  - Describes device's memory-access capabilities with on-chip memories

- CD & CDW Precisions (5 in all)

- Bit-Level, 16-bit Integer, 32-bit Integer, Single-Precision Floating-Point (SPFP), and Double-Precision Floating-Point (DPFP)

- IMB

- Block-based vs. Cache-based systems

## Devices Studied (18)

130 nm FMC	Ambric Am2045 <sup>1</sup>
	ClearSpeed CSX600
	Freescale MPC7447
90 nm RMC	Altera Stratix-II EP2S180
	ElementCXI ECA-64
	Mathstar Arrix FPOA
	Raytheon MONARCH
	Tilera TILE64
	Xilinx Virtex-4 LX200
	Xilinx Virtex-4 SX55
90 nm FMC	Freescale MPC8640D
	IBM Cell BE
65 nm RMC	Altera Stratix-III EP3SL340
	Altera Stratix-III EP3SE260
	Xilinx Virtex-5 LX330T
	Xilinx Virtex-5 SX95T
45 nm FMC	Intel Atom N270 <sup>2</sup>
40 nm RMC	Altera Stratix-IV EP4SE530

<sup>1</sup> Preliminary results based on limited vendor data (Ambric)

<sup>2</sup> Limited Atom cache data, not included in IMB results

# Metric Methodology

- CD for FPGAs

- Bit-level

- $f_{max}$  is max device frequency,  $N_{LUT}$  is number of look-up tables,  $W_i$  &  $N_i$  are width & number of fixed resources

$$CD_{bit} = f_{max} \times \left[ N_{LUT} + \sum_i W_i \times N_i \right]$$

- Integer

- Use method on right with Integer cores

$$CD_{int/FP} = (Ops_{DSP} + Ops_{LOGIC}) \times f_{achievable}$$

- Floating-point

- Use method on right with FP cores

Overhead - Reserve 15% logic resources for steering logic and memory or I/O interfacing

Memory-sustainable CD – Limit CD based on # of parallel paths to on-chip memory; each operation requires 2 memory locations

Parallel Operations – scales up to max. # of adds and mults (# of adds = # of mults)

Achievable Frequency – Lowest frequency after PAR of DSP & logic-only implementations of add & mult computational cores

IP Cores – Use IP cores provided by vendor for better productivity

## Integer & Floating-Point Analysis

Determine maximum amount of logic resources & maximum amount of special on-chip resources (e.g. DSP multipliers), for device



Determine resource utilization & maximum achievable frequency for one instance of core using DSP resources; repeat using logic-only resources



Determine number of simultaneous cores,  $Ops_{DSP}$ , that can be instantiated until all DSP resources are exhausted; repeat for logic-only resources to determine  $Ops_{LOGIC}$



Achievable frequency  $f_{achievable}$  is lower of frequencies determined in step 2 above



Iterate through combinations of DSP and logic-only cores to find an equal balance of addition and multiplication operations



# Metric Methodology

- CD for FMC and coarse-grained RMC devices

- Bit-level
- Integer
- Floating-point

$$CD_{bit} = f \times \left[ \sum_i W_i \times N_i \right]$$

$$CD_{int/FP} = f \times \sum_i \frac{N_i}{CPI_i}$$

- CDW for all devices

- Calculated using CD for each level of parallelism and dividing by power consumption at that level of parallelism
- CDW is *critical* metric for HPEC systems

$W_i$  - width of element type  $i$

$N_i$  - # of elements of type  $i$ , or # of instructions that can be issued simultaneously

$f$  - clock frequency

$CPI_i$  - cycles per instruction for element  $i$

## For all RMC

- Power scales linearly with resource utilization

## For FPGAs

- Vendor tools (PowerPlay, Xpower) used to estimate power for maximum LUT, FF, block memory, and DSP utilization at maximum freq.
- Maximum power is scaled by ratio of achievable frequency to maximum freq.

## For all FMC

- Use fixed, maximum power from vendor documentation

# Metric Methodology

- Internal Memory Bandwidth (IMB)

- Overall application performance may be limited by memory system

- Cache-based systems (CBS)

- Separate metrics for each level of cache
- Calculate bandwidth over range of hit rates

$$IMB_{cache} = \%hitrate \times \sum_i \frac{N_i \times P_i \times W_i \times f_i}{8 \times CPA_i}$$

- Block-based systems (BBS)

- Calculate bandwidth over a range of achievable frequencies
- For fixed-frequency devices, IMB is constant
- Assume most parallel configuration (wide & shallow configuration of blocks)
- Use dual-port configuration when available

$$IMB_{block} = \sum_i \frac{N_i \times P_i \times W_i \times f_i}{8 \times CPA_i}$$

*%hitrate* - Hit-rate scale factor

$N_i$  - # of blocks of element  $i$

$P_i$  - # of ports or simultaneous accesses supported by element  $i$

$W_i$  - width of datapath

$f_i$  - memory operating frequency, variable for FPGAs

$CPA_i$  - # of clock cycles per memory access

# Computational Density

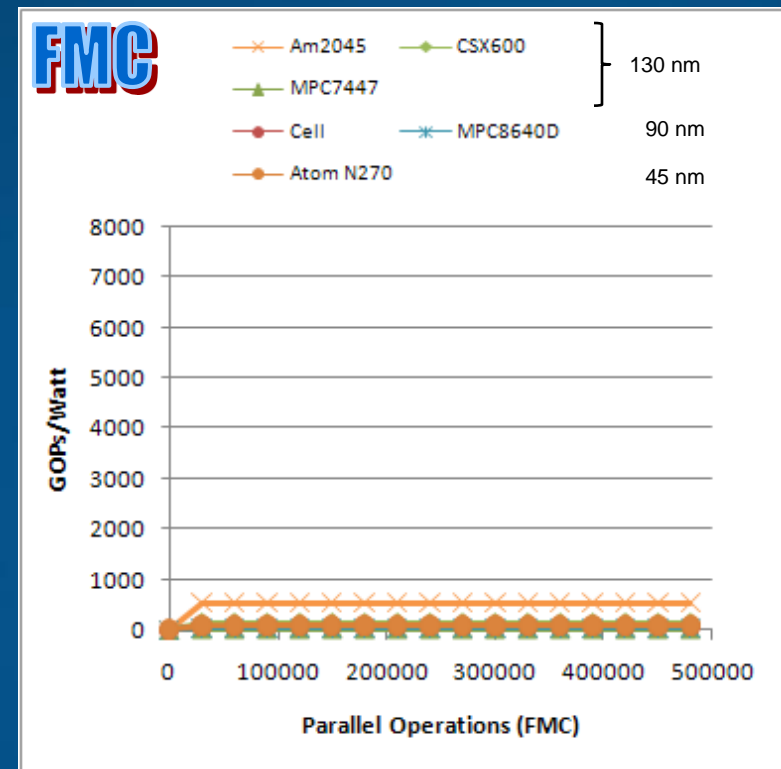
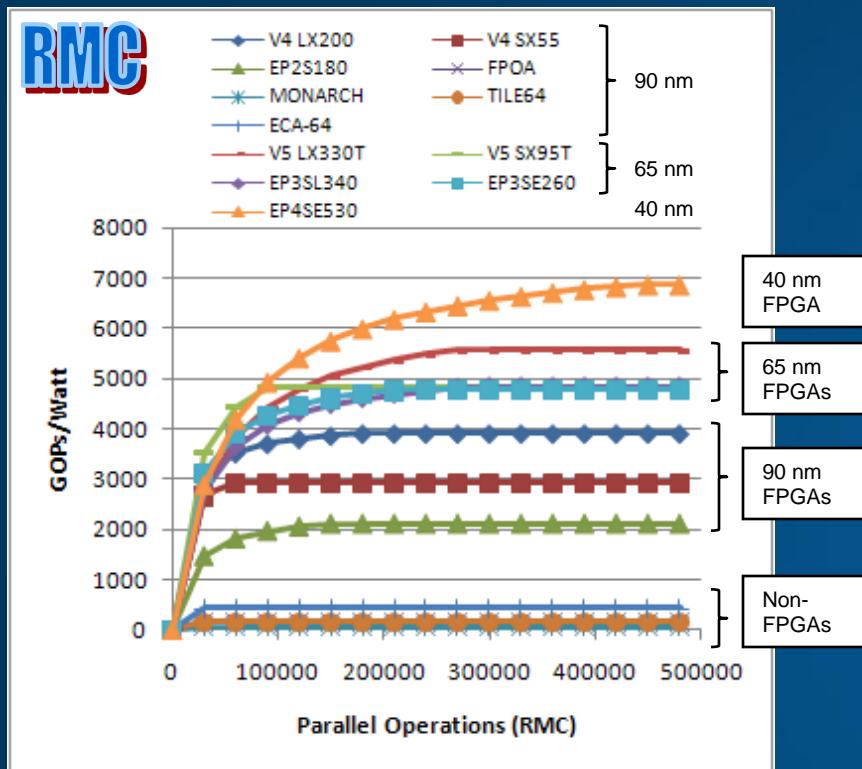
Device	Bit-level		16-bit Int.		32-bit Int.		SPFP		DPFP	
	Raw	Sustain.	Raw	Sustain.	Raw	Sustain.	Raw	Sustain.	Raw	Sustain.
Arrix FPOA	6144	6144	384	384	192	192				
ECA-64	2176	2176	13	13	6	6				
MONARCH	2048	2048	65	65	65	65	65	65		
Stratix-II S180	63181	63181	442	442	123	123	53	53	11	11
Stratix-III SL340	154422	154422	933	918	213	213	96	96	26	26
Stratix-III SE260	119539	119539	817	778	204	204	73	73	22	22
Stratix-IV SE530	243866	243866	990	766	312	312	171	171	88	88
TILE64	4608	4608	240	240	144	144				
Virtex-4 LX200	89952	89952	357	116	66	42	68	46	16	16
Virtex-4 SX55	29184	29184	365	110	71	40	31	31	7	7
Virtex-5 LX330T	150163	150163	606	300	131	122	119	116	26	26
Virtex-5 SX95T	48435	48435	599	226	221	92	82	82	15	15
Am2045	8064	8064	504	504	252	252				
Atom N270	307	307	14	14	8	8	8	8	5	5
Cell BE	4096	4096	205	205	115	115	205	205	19	19
CSX600	1536	1536	24	24	24	24	24	24	24	24
MPC7447	288	288	17	17	9	9	6	6	3	3
MPC8640D	576	576	34	34	18	18	12	12	6	6

RMC

FMC

- *Maximum* memory-sustainable CD is shown above (in GOPs)
- CD scales with parallel operations
- Various devices may have their highest CDs at different levels of parallelism
- Top CD performers are highlighted
- RMC devices perform best for bit-level & integer ops, FMC for floating-point
- Memory-sustainability issues seen when many, small registers are needed

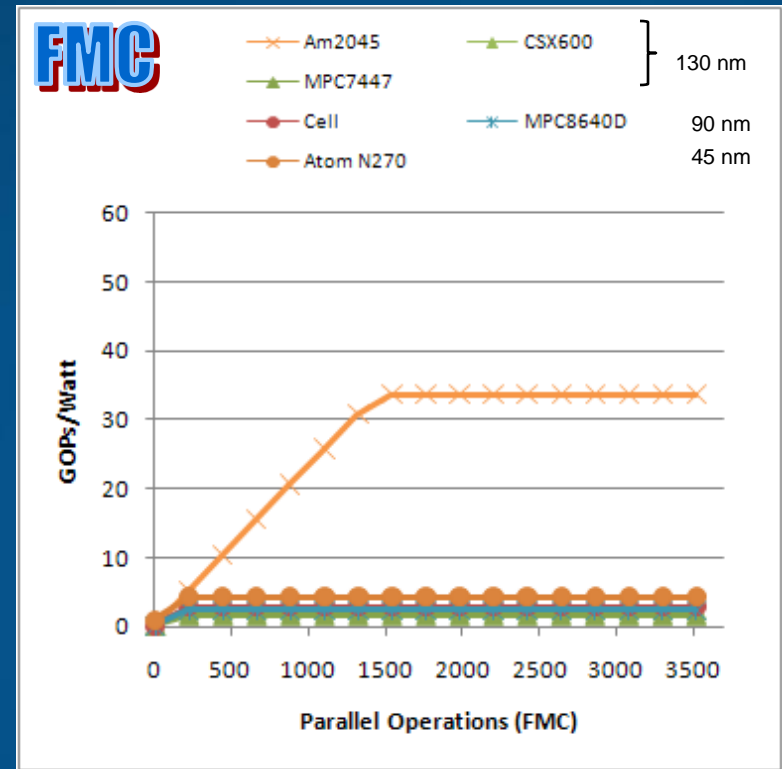
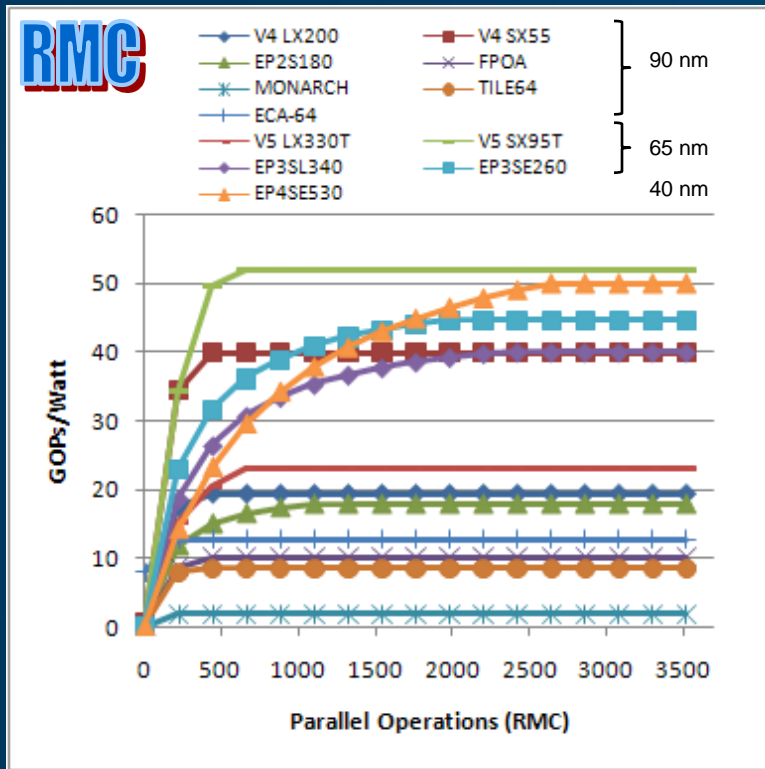
# Bit-level CDW



- RMC devices (specifically FPGAs) far outperform FMC devices
  - High bit-level CD due to fine-grained, LUT-based architecture
  - Low power
  - Power scaling with parallelism (area)

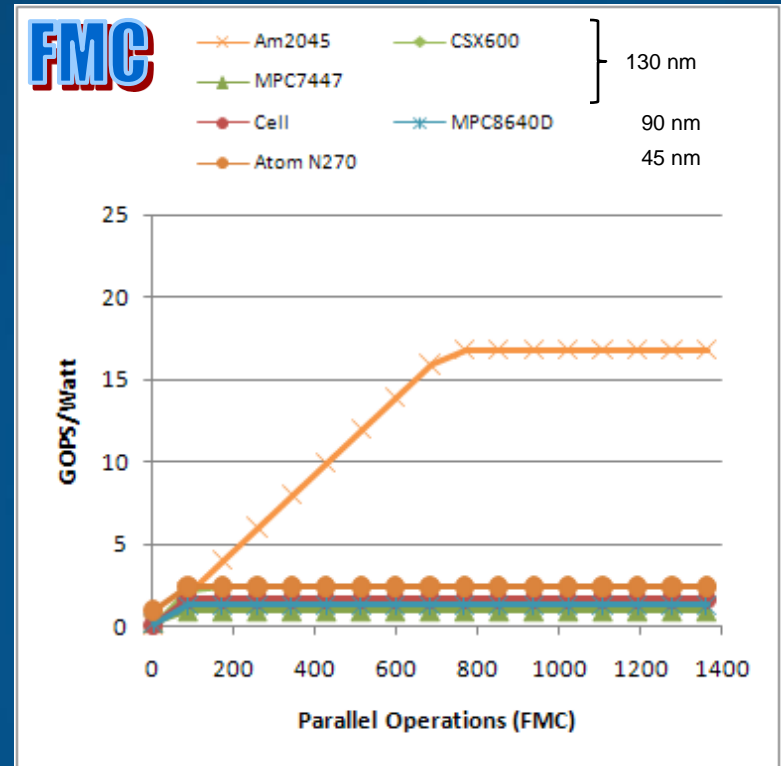
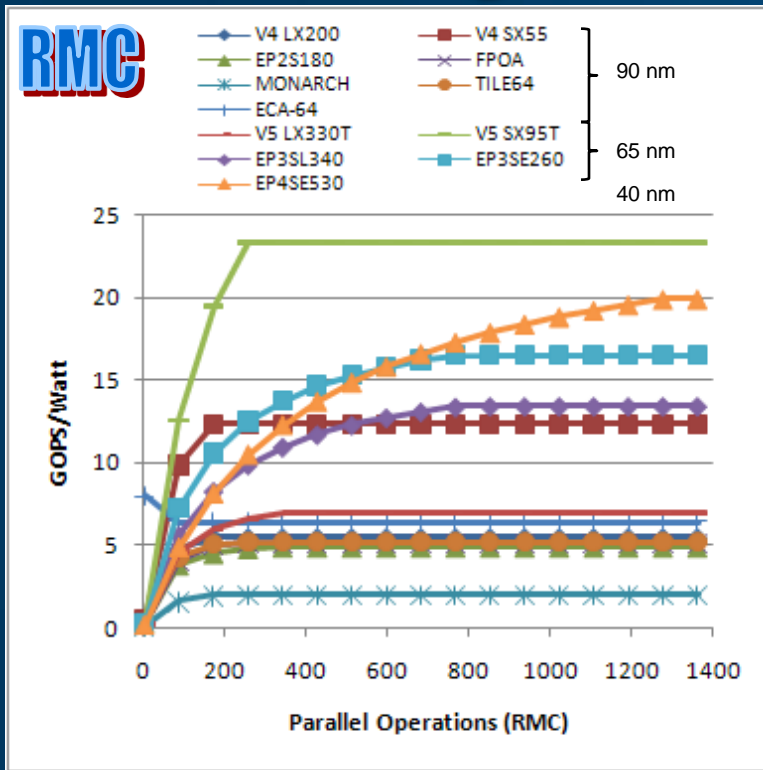
- **EP4SE530 (Stratix-IV) is best overall**
- 65 nm FPGAs are all strong performers
- V4 LX200 top performer of 90 nm devices
- Coarse-grained devices (both RMC & FMC) show poor performance

# 16-bit Integer CDW



- RMC devices outperform FMC
  - Low power
  - Power scaling with parallelism (area)
  - Requires algorithms that can take advantage of numerous parallel operations
  - Ambric (130 nm) shows promising prelim. results despite older process
- Virtex-4 SX55 is best performer in 90 nm class
- Strong performance from ECA-64 due to extremely low power consumption (one Watt at full utilization), despite low CD
- FPOA gives good, moderate performance due to high CD, but with higher power requirements
- **Virtex-5 SX95T (65 nm) is best overall with Stratix-IV EP4SE530 (40 nm) a close second**

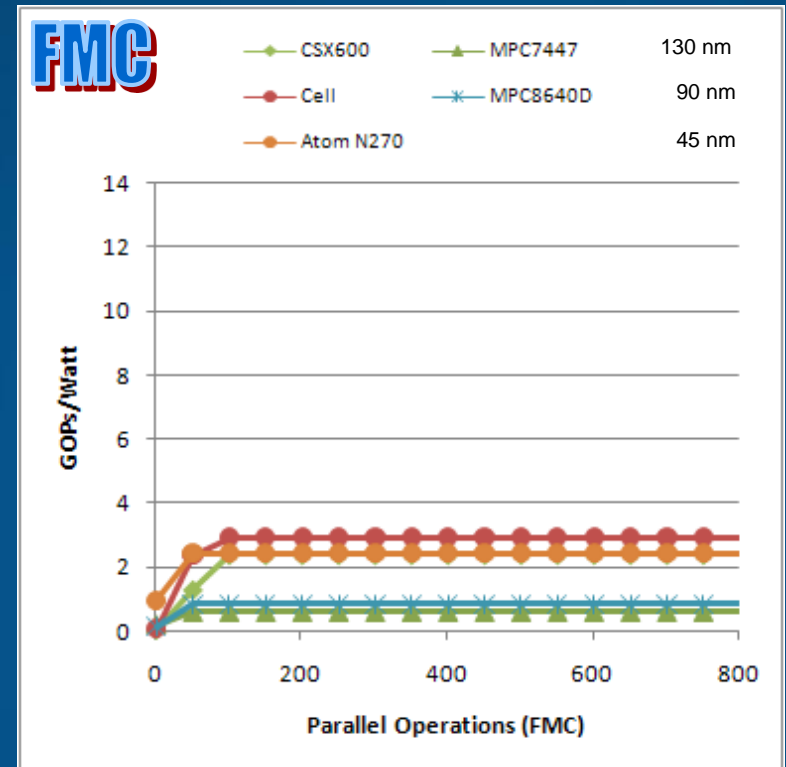
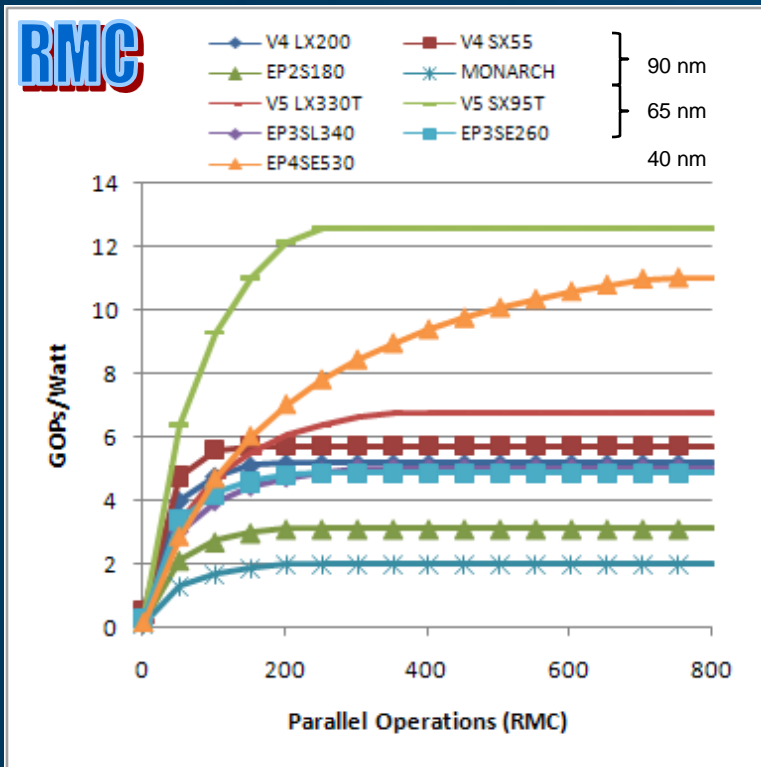
# 32-bit Integer CDW



- RMC devices outperform FMC
  - Low power
  - Power scaling with parallelism (area)
  - Requires algorithms that take advantage of numerous parallel operations
  - Ambric (130 nm) shows promising prelim. results despite older process

- For high levels of exploitable parallelism, the Virtex-4 SX55 is best in 90 nm class
- Strong performance from ECA-64 due to extremely low power consumption
- **Virtex-5 SX95T (65 m) is best overall**
- **SX devices benefit from low power consumption due to high DSP-to-logic ratio**

# SPFP CDW

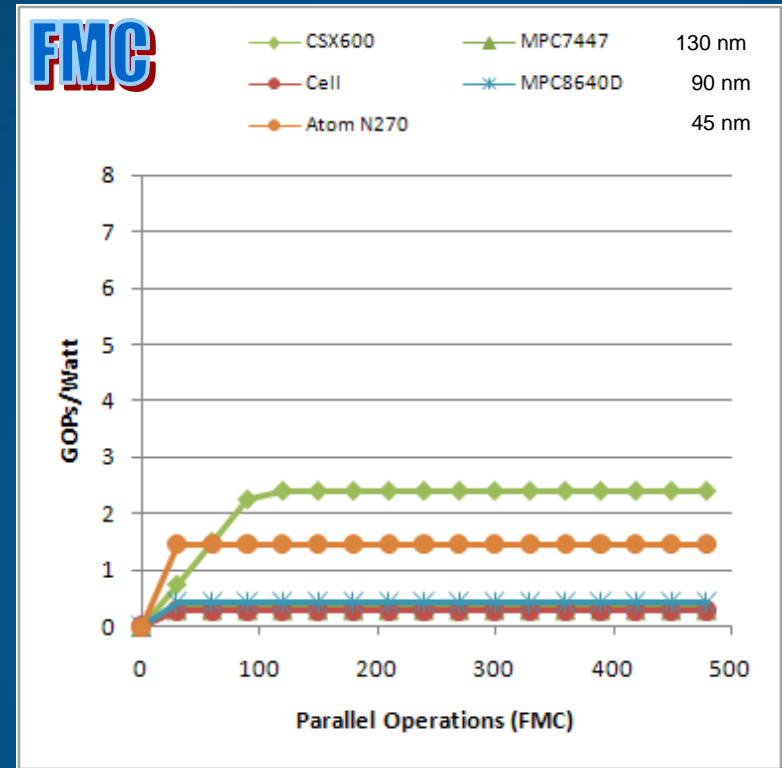
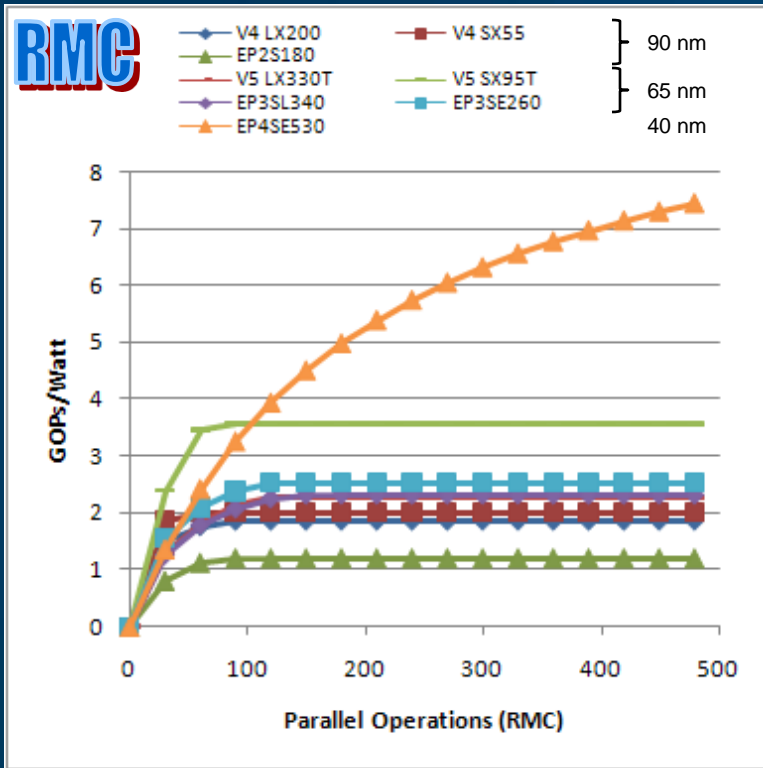


- RMC devices (specifically FPGAs) outperform FMC devices
  - Low power, especially FPGAs with large amount of DSP multiplier resources (consume less power than LUTs)
  - Power scaling with parallelism (area)
  - Devices not intended for floating-point computation (i.e. not designed to compete in current form) are excluded here (e.g. FPOA, TILE, ECA, Ambric)

- CSX600 modest due to average CD, low power
- Virtex-4 SX55 leads 90 nm due to power advantage
- Cell (90 nm) has large CD advantage, but very high power consumption hampers CDW capability
- **Virtex-5 SX95T (65 nm) has clear CDW advantage due to relatively high achievable frequency, high level of DSP resources, low power consumption of DSPs**

Note: we expect Altera FP CDW scores to improve when their new Floating-Point Compiler is used in place of current FP cores

# DPFP CDW



- RMC devices (specifically FPGAs) outperform most FMC devices

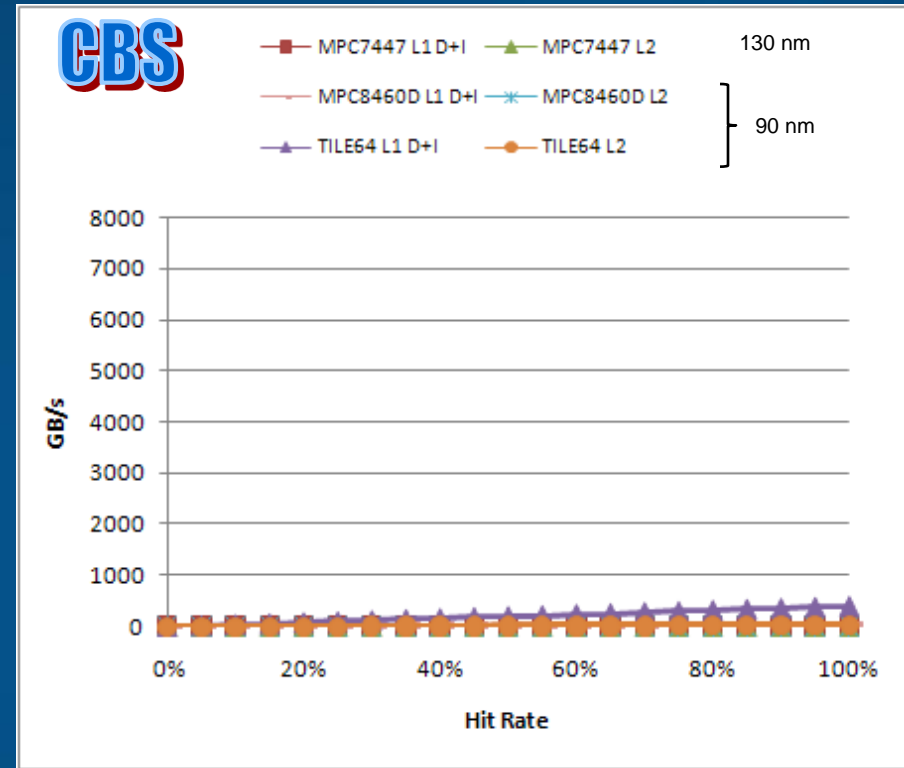
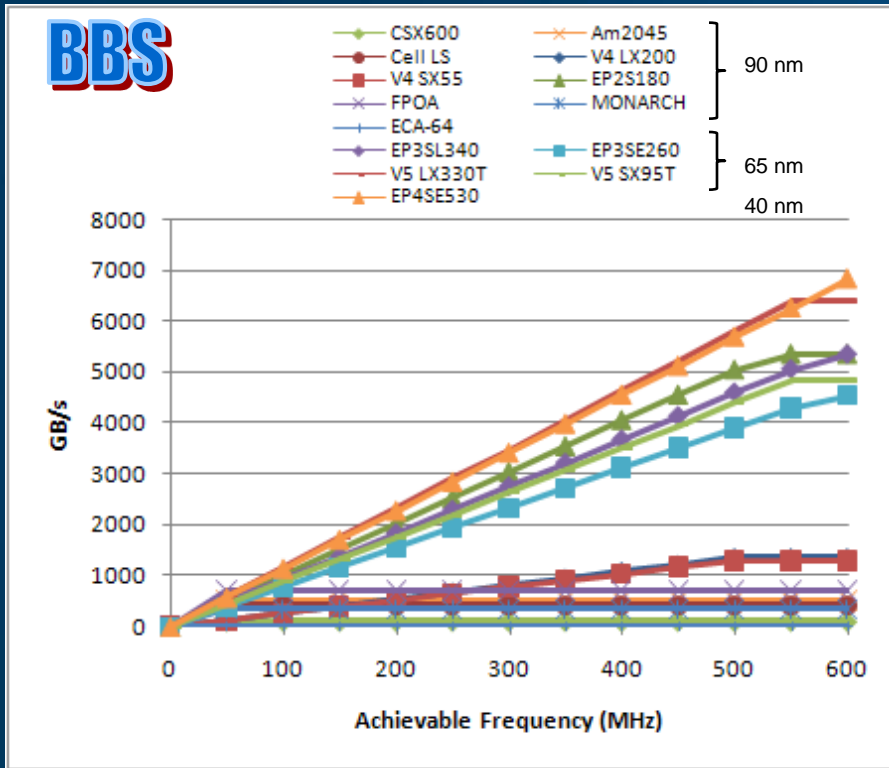
- Low power, especially FPGAs with large amount of DSP multiplier resources (consume less power than LUTs)
- Power scaling with parallelism (area)
- Devices not intended for floating-point computation are again excluded

- CSX600 (130 nm) performs better than several FPGAs due to high CD and moderate power
- SX devices (90 & 65 nm) perform well due to DSP power advantage, relatively high achievable frequencies
- **Stratix-IV EP4SE530 (40 nm) clear overall leader due to large fabric (DPFP cores are area-intensive)**

Note: we expect Altera FP CDW scores to improve when their new Floating-Point Compiler is used in place of current FP cores



# Internal Memory Bandwidth



- Block-based devices (specifically FPGAs) outperform cache-based devices
  - Many parallel paths to memory blocks
  - Can pack operands into wide data structures
  - Support for dual-port memories
  - Outperforms cache-based devices even on low frequency designs
  - IMB is constant for block-based fixed-frequency devices

- Cache-based systems (CBS)
  - MPC7447, MPC8640D perform poorly relative to most BBS devices
  - TILE64 (64 caches) does not compete with FPGAs
- Block-based systems (BBS)
  - **FPGAs dominate this metric**
  - Stratix-IV (40 nm) leads for higher-frequency designs, Virtex-5 leads for lower-frequency designs

# Future Work

- Compare algorithms using Computational Intensity (CI) metric

$$CI = \frac{\text{Arithmetic Operations}}{\text{Memory Operations}}$$

- Use CD, IMB, and CI metrics to correlate device characteristics and application characteristics

```

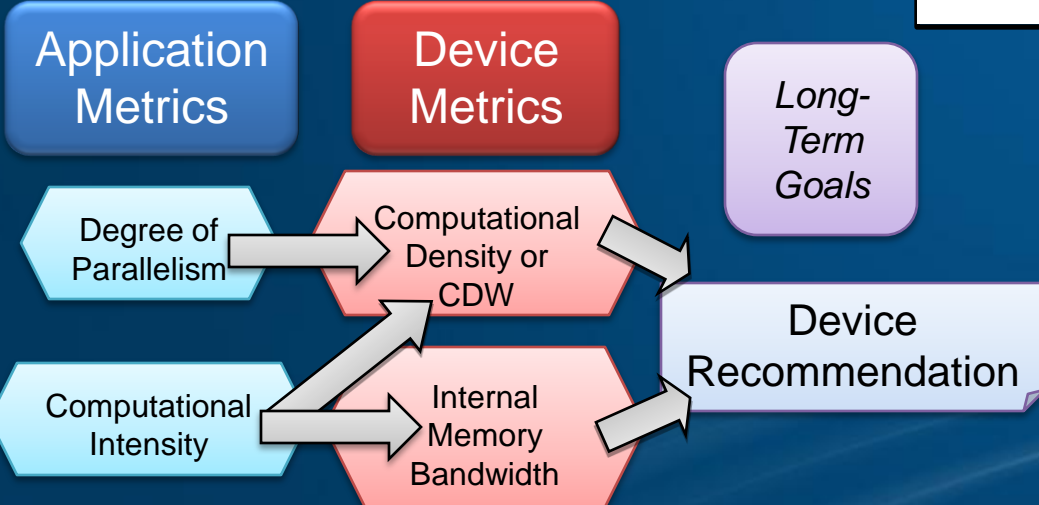
size2 = floor(size/2);
% For each pixel in the image
for i = 1:size2
    for j = 1:size2
        % clear the window sum
        accum_win = 0;
        % clear the number of pixels averaged
        num_denom = 0;
        % For each pixel in the window
        for i2 = -size2:size2
            win_i = i + i2;
            if (win_i > 0 && win_i < 513)
                for j2 = -size2:size2
                    win_j = j + j2;
                    if (win_j > 0 && win_j < 513)
                        % increase number of elements
                        % added to window
                        num_denom = num_denom + 1;
                        % gather window sum
                        accum_win = uint32(accum_win +
                            uint32(noisy(win_i, win_j)));
                    end
                end
            end
        end
        % perform filter
        cln_img(i, j) = uint8(accum_win / num_denom);
    end
end
    
```

$Z * I = M$   
Overall Arithmetic Operations:  
 $I * M$

Pixel Level Loop:  
 $8s^2 + 7s = Z$

Innermost loop:  $8s$

Mem Operations:  
 $s^2 * I^2$



**2D-Convolution (I = Image size and s = filter size)**  
 For I = 512; s = 3; Computational Intensity = 9.9  
 For I = 512; s = 7; Computational Intensity = 8.9  
 For I = 512; s = 15; Computational Intensity = 8.5  
**CFAR** - Computational Intensity = 2.1  
**Radix-4 FFT** - Computational Intensity = 4.7  
**Direct Form FIR** - Computational Intensity = 4.1  
**Matrix Multiply** - Computational Intensity = 2.0

# Summary

	Best Overall	Best RMC	Best FMC	Best of 90 nm & larger proc.
Bit-level CDW	<b>EP4SE530</b>	EP4SE530	Am2045	V4 LX200
16-bit Integer CDW	<b>V5 SX95T</b>	V5 SX95T	Am2045	V4 SX55
32-bit Integer CDW	<b>V5 SX95T</b>	V5 SX95T	Am2045	V4 SX55
SPFP CDW	<b>V5 SX95T</b>	V5 SX95T	Cell	V4 SX55
DPFP CDW	<b>EP4SE530</b>	EP4SE530	CSX600	CSX600
IMB	<b>EP4SE530</b>	EP4SE530	Am2045	EP2S180



# Conclusions

- RC Taxonomy & Reconfigurability Factors
  - Provides framework for comparing RMC & FMC devices
  - Develops concepts and terminology to define characteristics of various computing device technologies
- CD and CDW Metrics
  - Basis to compare devices on computational performance & power
    - Large variations in resulting data when applied across disparate device suite
    - **FPGAs with many low-power DSPs** tended to have very high CDW scores, even for single-precision, floating-point operations
  - With increasing importance of energy, **CDW** becomes a critical metric
- IMB Metric
  - Basis to compare devices for on-chip memory access capabilities
  - Block-based systems tended to outperform cache-based systems
- Architecture reformation & Moore's law
  - Explicit parallelism allows for full utilization of process technology & transistor count improvements



# Acknowledgements

This work was made possible by

- NSF I/UCRC Program (Center Grant EEC-0642422)
- CHREC members (31 industry & govt. partners)
  
- Altera Corporation (equipment, tools)
- MathStar Incorporated (equipment, tools)
- Xilinx Incorporated (equipment, tools)

## Questions?



# References

- Altera Corp., *Stratix II Device Handbook*, 2007.
- Altera Corp., *Stratix III Device Handbook*, 2007.
- Altera Corp., *Stratix IV Device Handbook*, 2008.
- Ambric, Inc., “Technology Overview,” <http://www.ambric.com/technology/technology-overview.php>.
- M. Barton, “Tilera’s Cores Communicate Better,” *Microprocessor Report*, Nov. 2007.
- T. Chen, et al., “Cell Broadband Engine Architecture and its First Implementation--A Performance View,” *IBM Journal of Research & Development*, vol. 51, no. 5, Sept. 2007, pp. 559-572.
- ClearSpeed Technology PLC, *CSX600 Architecture Whitepaper*, 2007.
- A. DeHon. Reconfigurable Architectures for General Purpose Computing, PhD thesis, MIT AI Lab, Sept. 1996.
- Element CXI, Inc., *ECA-64 Device Architecture Overview*, 2007.
- Element CXI, Inc., *ECA-64 Product Brief*, 2007.
- Freescale Semiconductor, Inc., *AltiVec Technology Programming Environments Manual Rev. 3*, 2006.
- Freescale Semiconductor, Inc., *MPC7450 RISC Microprocessor Family Reference Manual Rev. 5*, 2005.
- Freescale Semiconductor, Inc., *MPC8641D Integrated Host Processor Family Reference Manual Rev. 2*, 2008.
- T. Halfhill “Ambric’s New Parallel Processor,” *Microprocessor Report*, Oct. 2006.
- Intel Corp., *Intel 64 and IA-32 Architectures Software Developer’s Manual Volume 1: Basic Architecture*, Apr. 2008.
- Intel Corp., *Mobile Intel Atom Processor N270 Single Core Datasheet*, May 2008.
- Mathstar, Inc., *Arrix Family FPOA Architecture Guide*, 2007.
- Mathstar, Inc., *Arrix Family Product Data Sheet & Design Guide*, 2007.
- Raytheon Company, *World’s First Polymorphic Computer - MONARCH*, 2006.
- D. Strenski, “FPGA Floating Point Performance -- a pencil and paper evaluation,” *HPCWire*, Jan. 12, 2007, <http://www.hpcwire.com/hpc/1195762.html>.
- Tilera Corp., *TILE64 Processor Product Brief*, 2008.
- D. Wang, “ISSCC 2005: the Cell Microprocessor,” *Real World Technologies*, Feb. 2005, retrieved Jan. 2008, <http://www.realworldtech.com/page.cfm?ArticleID=rwt021005084318&p=2>.
- Xilinx, Inc., *Virtex-4 Family Overview*, 2007.
- Xilinx, Inc., *Virtex-5 Family Overview*, 2008.

The background is a solid blue color. In the upper left quadrant, there are several thin, white, curved lines that sweep across the top of the page, creating a sense of motion or flow.

**BACKUP**

# Devices Studied

## FMC Device Features

	Device	Cores	Instructions Issued/Core	Datapath Width (bits)	Frequency (MHz)	Power (W)	On-chip Memory
130 nm	Am2045	360	3+1	32	350	15	45 bricks ea. w/ 8 SRAM banks
	CSX600	1+96	1	64	250	10	I, D caches, <b>96 32-bit banks SRAM</b>
	MPC7447	1+1	1+2 Int, 2+1 SPFP, 3 DPFP	32/128	1000	10	L1-I, L1-D: 4 words/access @ 2 cycles/access, L2: 8 words/access @ 9 cycles/access
90 nm	Cell BE	1+8	2+1	64/128	3200	70	L1-I, L1-D, L2 (PPE), <b>8 128-bit LS banks (SPEs)</b>
	MPC8640D	2+2	1+2 Int, 2+1 SPFP, 3 DPFP	32/128	1000	14	Ea. core: L1-I, L1-D: 4 words/access @ 2 cycles/access, L2: 8 words/access @ 11.5 cycles/access
45 nm	Atom N270	1+1	1+1	64/128	1600	3.3	Unknown

## FPGA Device Features

	Device	LUTs	DSPs	Max. Frequency (MHz)	Min. Power (W)	Max. Power (W)	On-chip Memory
90 nm	Stratix-II EP2S180	143,520	768	500	3.26	30	9 128-bit dual port blocks @ 420 MHz, 768 32-bit dual port blocks @ 550 MHz, 930 16-bit dual port blocks @ 500 MHz
	Virtex-4 SX55	49,152	512	500	1	10	48 72-bit dual port blocks @ 600 MHz, 864 32-bit dual port blocks @ 580 MHz,
	Virtex-4 LX200	178,176	96	500	1.27	23	48 72-bit dual port blocks @ 600 MHz, 1040 32-bit dual port blocks @ 580 MHz,
65 nm	Stratix-III EP3SE260	203,520	768	550	2.11	25	320 32-bit dual port blocks @ 500 MHz
	Stratix-III EP3SL340	270,400	576	550	2.83	32	336 32-bit dual port blocks @ 500 MHz
	Virtex-5 SX95T	58,800	640	550	1.89	10	488 72-bit dual port blocks @ 550 MHz
	Virtex-5 LX330T	207,360	192	550	3.43	27	648 72-bit dual port blocks @ 550 MHz
40 nm	Stratix-IV EP4SE530	424,960	1,024	600	3.55	39	64 72-bit dual port blocks @ 600 MHz, 1280 32-bit dual port blocks @ 600 MHz,



# Devices Studied

## Other RMC Device Features

	Device	PE	Frequency (MHz)	Min. Power (W)	Max. Power (W)	On-chip Memory
90 nm RMC	ElementCXI ECA-64	64 16-bit hetero. elements	200	0.05	1	4 16-bit memory units, 5 simultaneous operations
	Mathstar Arrix FPOA	256 16-bit ALUs, 64 16x16 MACs	1000	18.82 @ 25%	46.25 @ 100%	80 32-bit dual port banks @ 1 GHz, 12 72-bit single port banks @ 500 MHz
	Raytheon MONARCH	6 32-bit RISC processor cores, 12 256-bit Arithmetic Clusters	333	6.7	33	31 memory clusters, 4 memories/cluster, dual ported, 32 bits wide
	Tilera TILE64	64 32-bit 3 issue VLIW processor cores	750	5.11	28	64 32-bit L1 I, D caches, Unified L2 cache @ 7 cycle access

## FPGA Achievable Frequencies

Device	Bit-Op	16-bit Int.	32-bit Int.	SPFP	DPFP
Stratix-II EP2S180	500	420	410	286	148
Stratix-III EP3SE260	550	273	400	329	195
Stratix-III EP3SL340	550	273	400	329	195
Stratix-IV EP4SE530	550	243	291	241	184
Virtex-4 SX55	500	249	344	274	185
Virtex-4 LX200	500	249	344	274	185
Virtex-5 SX95T	550	378	463	357	237
Virtex-5 LX330T	550	378	463	357	237

Stratix-III &-IV Bit-Op frequency limited by max DSP frequency