



# **The Design and Applications of BEE2: A High End Reconfigurable Computing system**

*Chen Chang, John Wawrzynek,  
Bob Brodersen,  
EECS, University of California at Berkeley*

---



# High-End Reconfigurable Computer (HERC)

---

- A computer with *supercomputer-like performance*, based solely on FPGAs and/or other reconfigurable devices as the processing elements.
  - The inherent fine-grain flexibility of the FPGAs allow all data-paths, control, memory ports, and communication channels to be customized on a per-application basis and parallelism to be exploited at all levels.
- BEE2 development is underway:
  - demonstrate of the concept;
  - It will motivate engagement of application domain experts;
  - Motivate creative thinking in software/programming tools;
  - Be the first in a series of machines;
  - It is a joint project with Xilinx using their 130nm devices.
- Based on concepts demonstrated in BEE2 prototype, 1 petaOPS ( $10^{15}$ ) in 1 cubic meter attainable within 3 years.




# Applications Areas of Interest

---



- **High-performance DSP**
  - SETI Spectroscopy, ATA / SKA Image Formation
  - Hyper-spectral Image Processing (DARPA)
- **Scientific computation and simulation**
  - E & M simulation for antenna design (BWRC)
  - Fusion simulation (UW)
- **Communication systems development Platform**
  - Algorithms for SDR and Cognitive radio
  - Large wireless Ad-Hoc sensor networks
  - In-the-loop emulation of SOCs and Reconfigurable Architectures
- **Bioinformatics**
  - BLAST (Basic Local Alignment Search Tool) biosequence alignment
  - Molecular Dynamics (Drug discovery)
- **System design acceleration**
  - Full Chip Transistor-Level Circuit Simulation (Xilinx)
  - RAMP (Research Accelerator for MultiProcessing)



# Radio Astronomy (1MHz~500GHz)




## Electromagnetic Spectrum



Colliding black holes

August 16th, 2005




Radio jets

Image courtesy of NRAO/AUI



# Allen Telescope Array (CA, USA)




Artist's conception  
Hat Creek, California  
350 Antenna



# Allen Telescope Array (cont.)



0.5~11GHz RF



August 16th, 2005

EECS, UC Berkeley

6




# Large-N, Small-D Concept

---

- Use lots of small diameter antennas to achieve large aggregate collecting area
- Benefits
  - Extremely high quality coverage
  - Very wide range of baseline lengths
  - Flexible usage model, multi-user, multi-subarrays
  - Reliability through redundancy
  - Economy of scale



# Telescope cost for fixed collect area







# Problems with existing approach

---

- All specialized instrument design
  - Separate PCB for each subsystem, dedicated functionality
  - Custom interconnect, backplane, and memory interface
  - Fully global synchronous I/O and processing
    - Clock distribution, power consumption, and voltage regulation
- Each instrument design cycle is 5 years!!!
- Instrument upgrade takes the similar effort as designing a new product



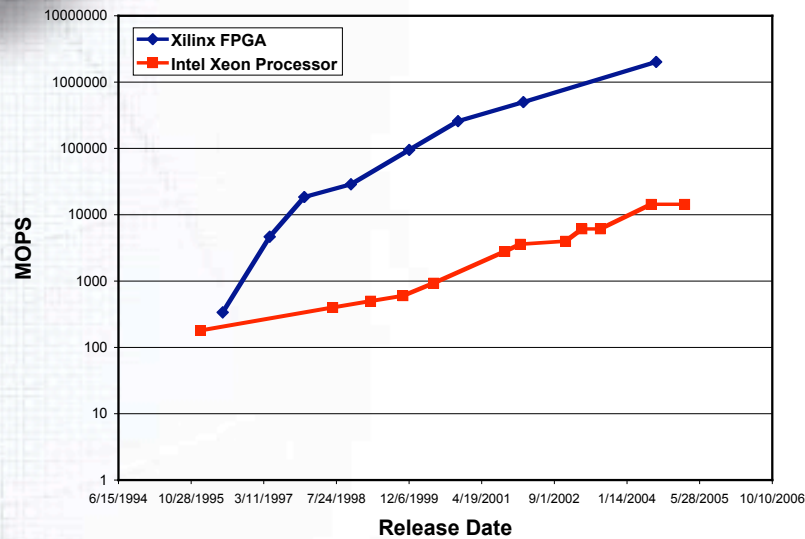
# Why not use...

---

- Microprocessor/DSP clusters?
  - Multi-processor programming is extremely hard, especially for real-time applications
  - Limited I/O capability, high power consumption, low computational density
- ASIC?
  - Lack of flexibility
  - Long design cycles
- FPGA? Sure!

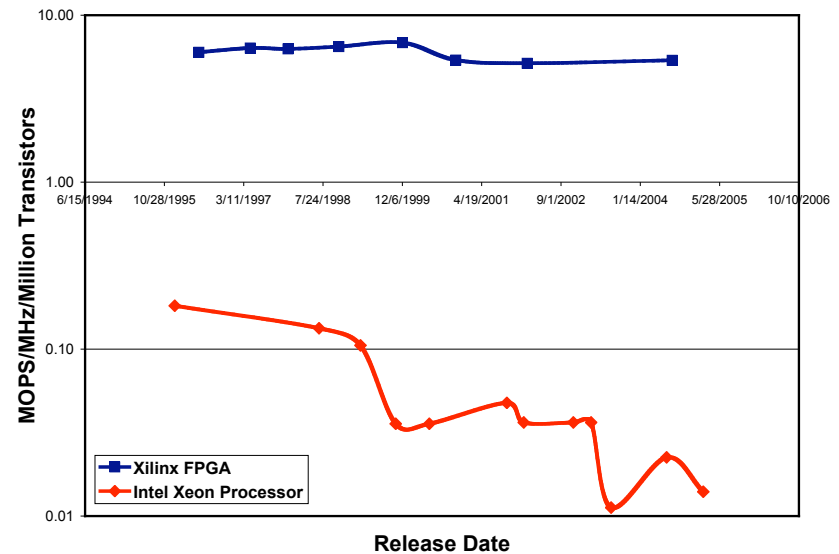


# Moore's Law in FPGA world



**100X higher performance  
100X more efficient  
than microprocessors!**

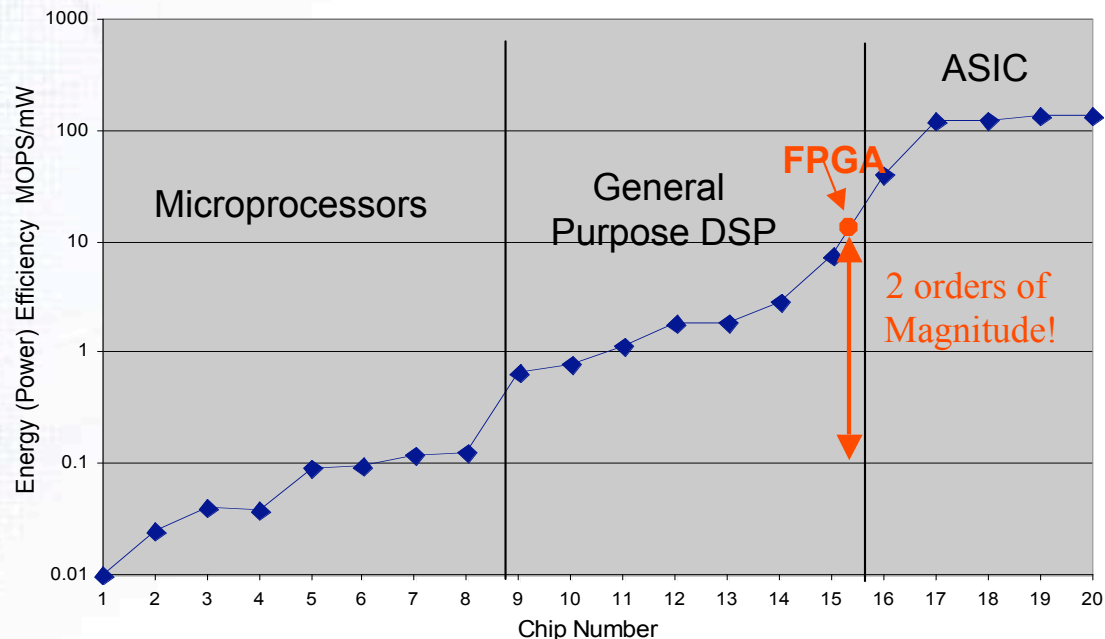
**FPGA performance  
doubles every 12 month!**





# Energy Efficiency (MOPS/mW)

- Based published results at ISSCC conferences and our measured results (FPGA).



- Specialized circuits use less energy per operation.
- Inherent computation density means devices can run at lower speed consuming less power.
- Reduced power consumption is a priority for FPGA vendors.




# BEE2 system design philosophy

---

- Compute-by-the-yard
  - Modular computing resource
  - Flexible interconnect architecture
  - On-demand reconfiguration of computing resources
- Economy-of-scale
  - Ride the semiconductor industry Moore's Law curve
  - All COTS components, no specialized hardware
  - Survival of application software using technology independent design flow



# BEE2 Compute Module



August 16th, 2005

EECS, UC Berkeley



14X17 inch 22 layer PCB

14



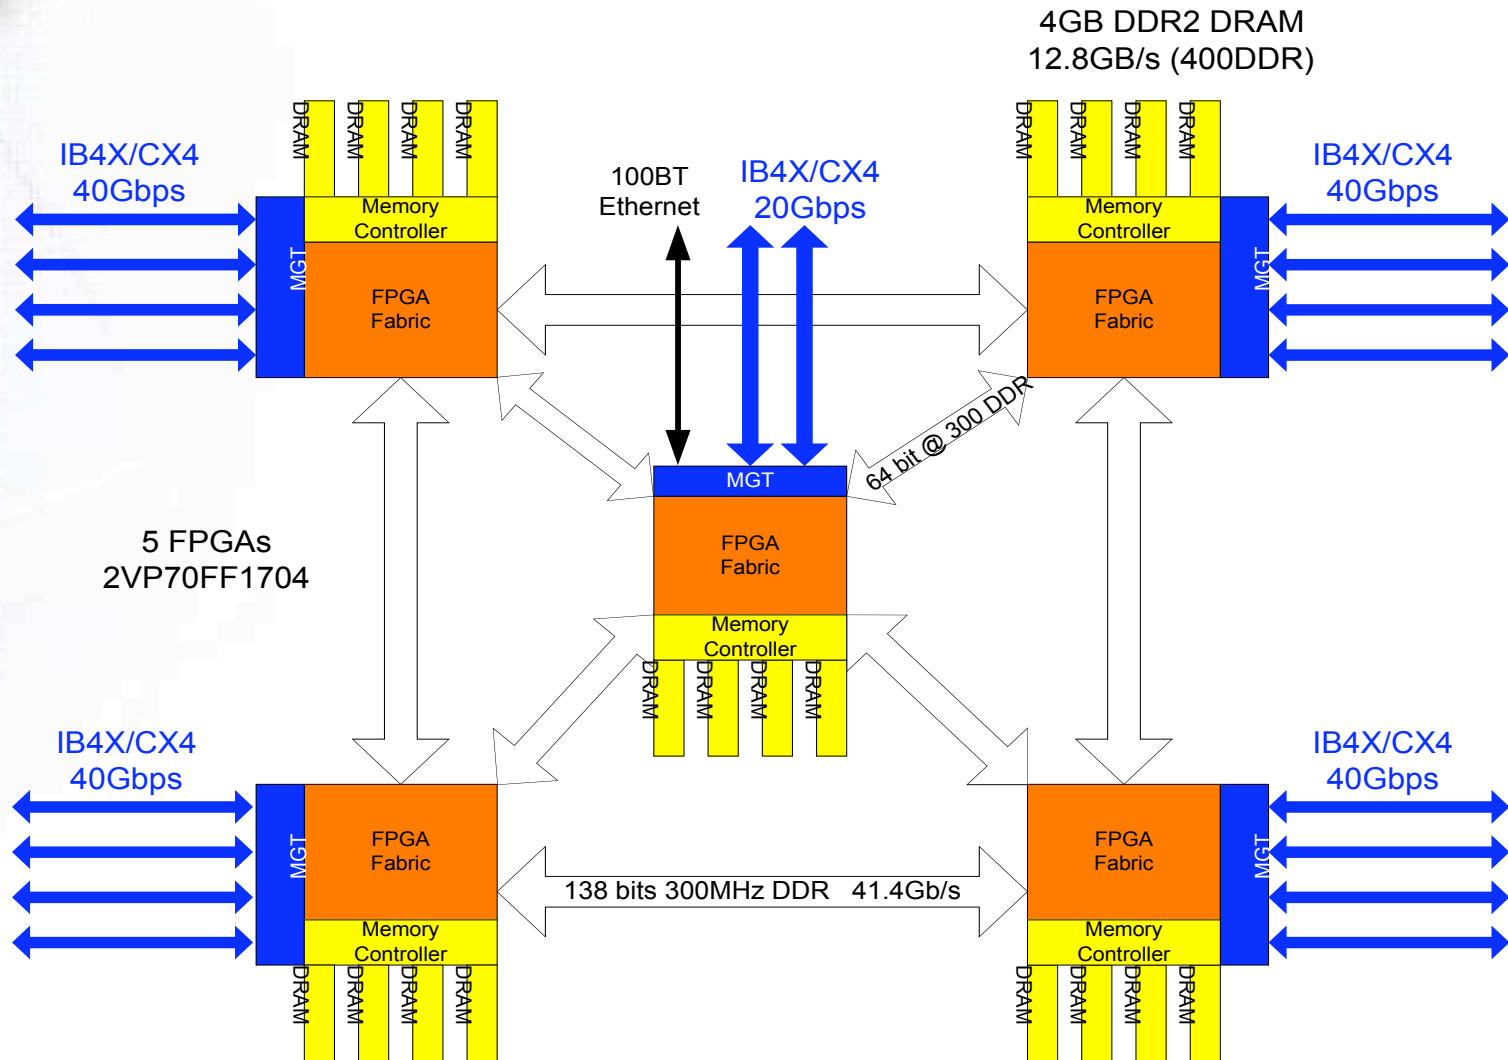
# Basic Computing Element

- Single Xilinx Virtex 2 Pro 70 FPGA
  - ~70K logic cells (35K slices)
  - 2 PowerPC405 cores
  - 324 dedicated multipliers (18-bit)
  - 5.7 Mbit SRAM on-chip
  - 20X 3.125-Gbit/s duplex serial communication links (MGTs)
- 4 physical DDR2-400 banks
  - Each banks has 72 data bits with ECC
  - Independently addressed with 32 banks total
  - Up to 12.8 GBps memory bandwidth, with maximum 4 GB capacity





# Compute Module Diagram





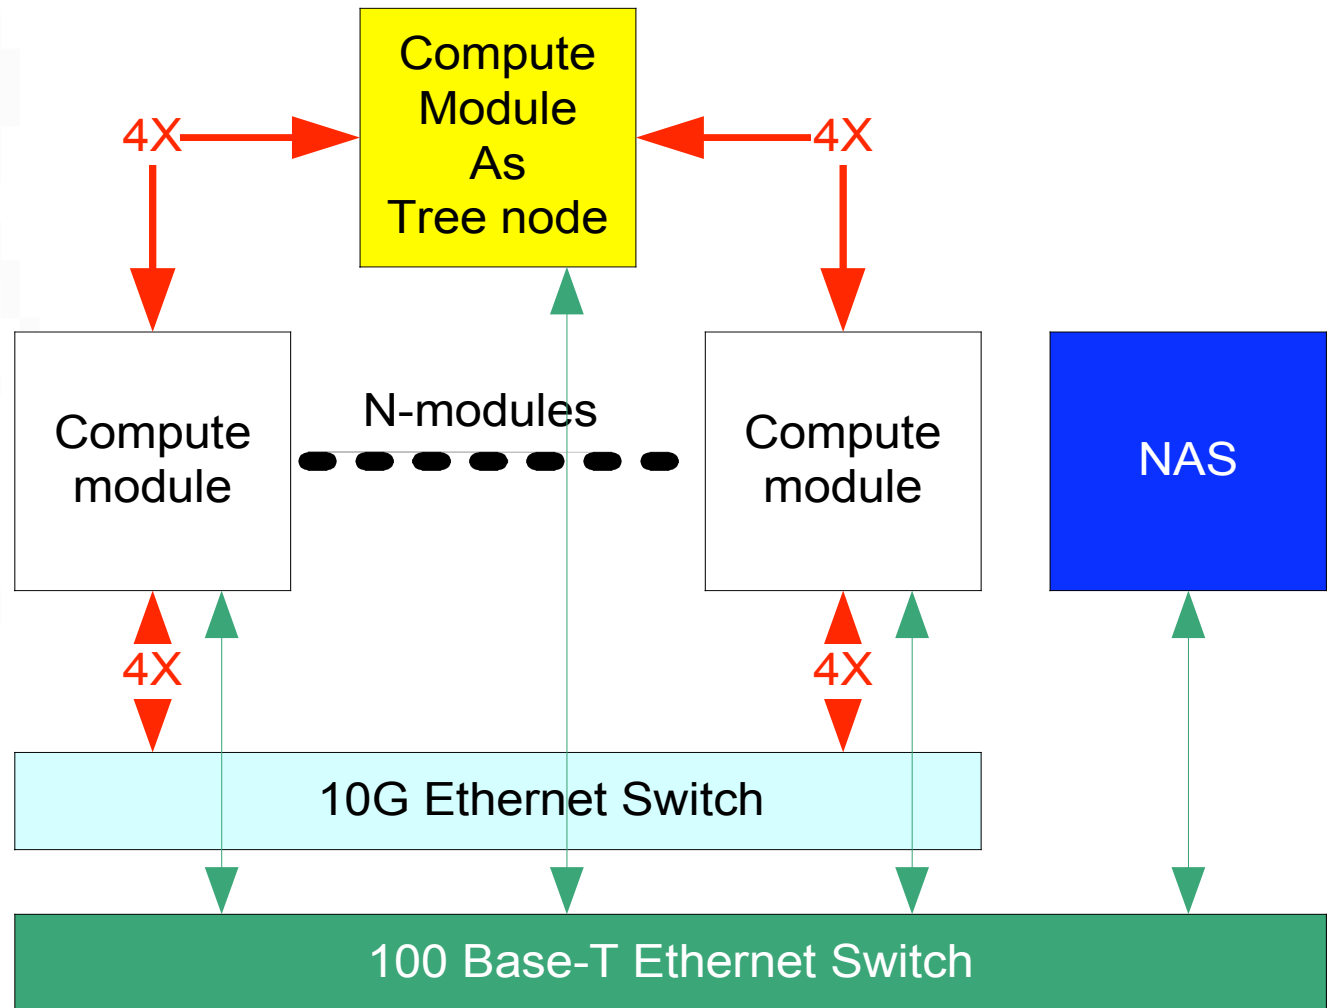


# Inter-Module Connections

Global  
Communication  
Tree

Stream Packets

Admin, UI, NFS





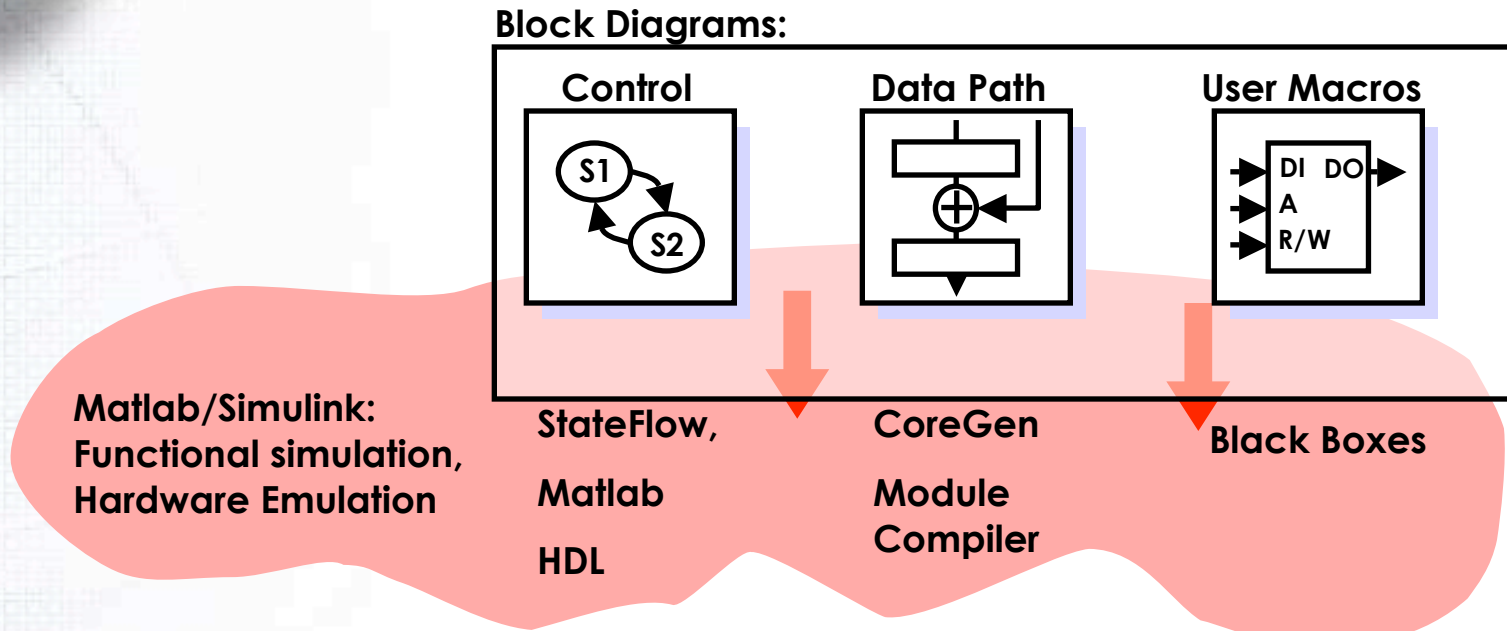
# 19" Rack Cabin Capacity

- 40 compute nodes in 5 chassis (8U) per rack
- ~40TeraOPS, ~1.5TeraFLOPS
- 150 Watt AC/DC power supply to each blade
- ~6 Kwatt power consumption
- Hardware cost: ~ \$500K





# Programming Model : Discrete Time Block Diagram with FSM



- Xilinx system generator library with BEE2 hardware specific hardware abstractions
- User assisted portioning with automatic system level routing



# BEE2 hardware abstractions

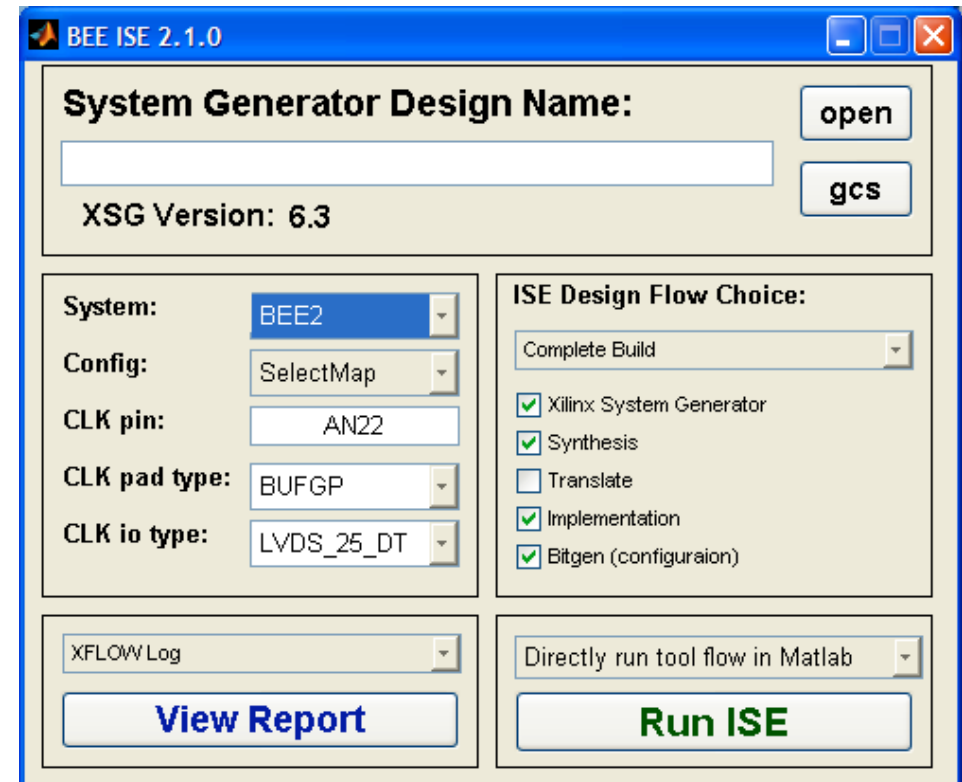
---

- Data flow operators
  - Data type: fix-point
  - Math operators: +/-, \*, /, &, |, xor, ~, >, =, <, srl, sll, sra
  - Control operators: demux/switch, mux/merge
- Memory
  - On-chip SRAM/Registers: shift register, RAM, ROM
  - Off-chip DRAM: stream RAM
- Communication and I/O
  - Static links: stream I/O
  - Dynamic links: Remote DMA
- Synchronization
  - Time stamp



# Tool flow

- Xilinx ISE 6.3i SP3
- System Generator 6.3i
- Synplify Pro 7.7.1
- Matlab 7 / Simulink 6
- BEE\_ISE 2.1.0
  - Tool flow wrapper in Matlab GUI
  - Automate CAD flow parameter optimizations and hardware system specific parameters
  - Require minimal knowledge of the tools flow from the end users





# Radio Astronomy Applications

Focused here first, because

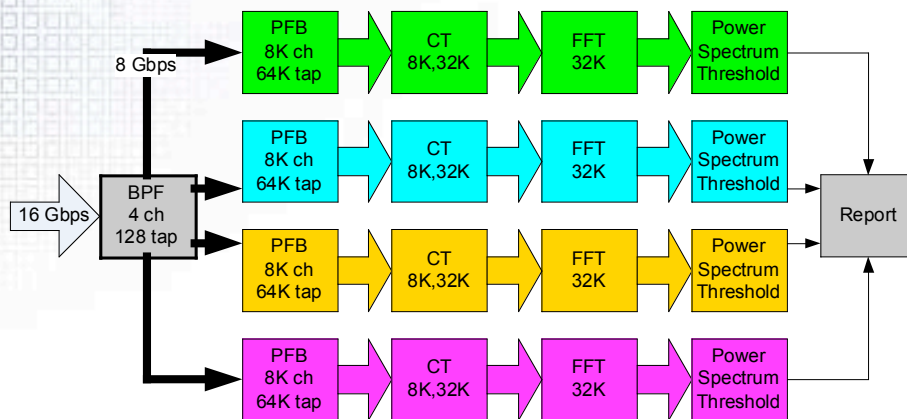
- Experienced with DSP-like problems
- App works well with our existing programming tools

- SETI Spectrometer

- Target: 0.7Hz channels over 800MHz → 1 billion Channel real-time spectrometer

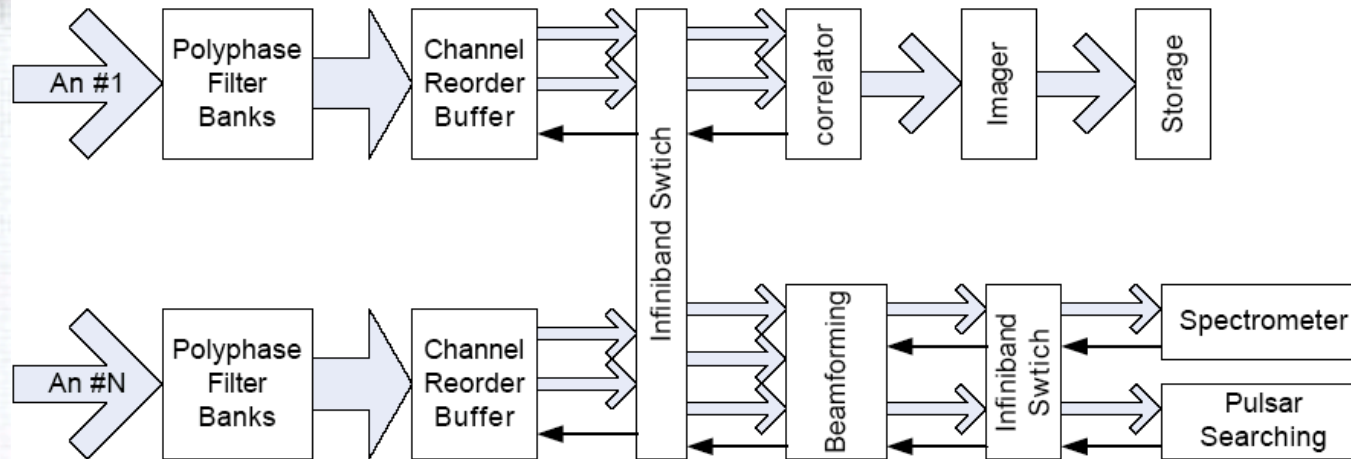
- Results:

- One BEE2 module meets target and yields 333GOPS (16-bit mults, 32-bit adds), at 150Watts (similar to desk-top computer)
    - >100x peak throughput of current Pentium-4 system on integer performance, & >100x better throughput per energy.





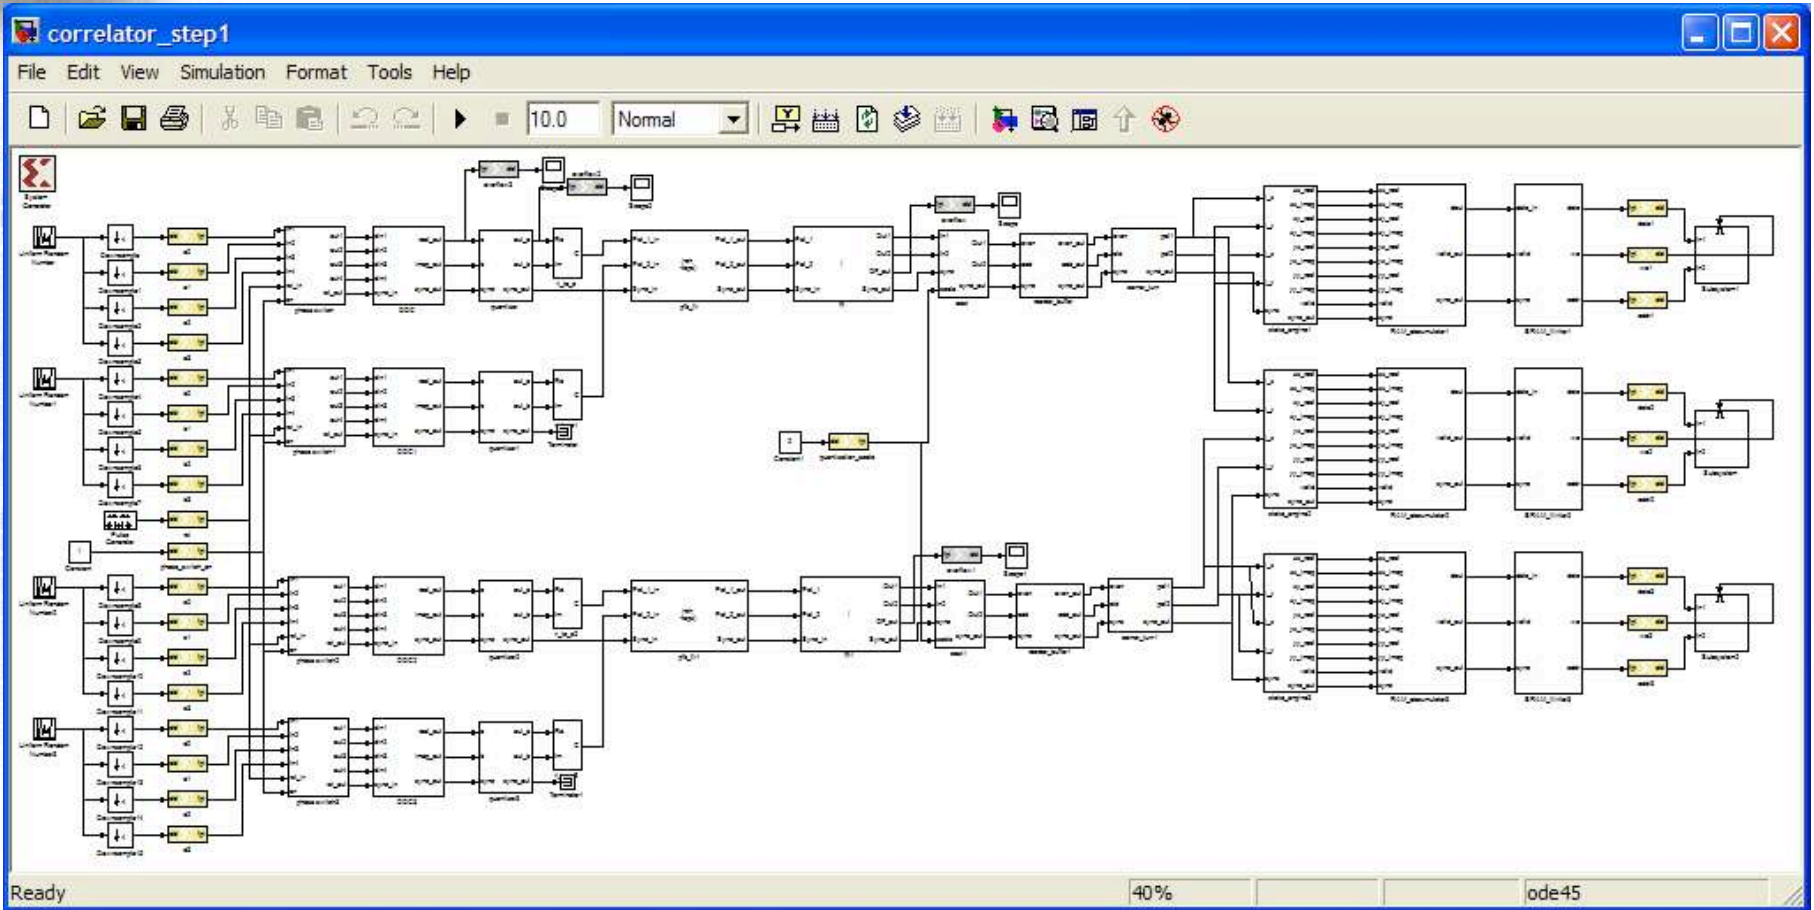
# Unified Radio Astronomy DSP processing architecture



Antenna array size	32	206	350	Notes:	100MHz band from each antenna
Baselines (# correlations)	496	21115	61075		
BEE2 modules (PFB)	1	7	11		Dual polarization
BEE2 modules (XMAC)	1	43	123		
BEE2 modules (Imager)	1	29	82		512x512 pixels image
Digitizers	4	26	44		1024 frequency channels per pixel
Ethernet switches	2	31	86		
<b>estimated \$ total(K)</b>	<b>\$ 99.20</b>	<b>\$ 2,036.00</b>	<b>\$ 5,441.00</b>		1 image dump / second



# 4 antenna correlator designs







# Benchmark applications


---

- 1024 channel dual polarization Polyphase Filter Bank (PFB) with 8K tap filter coefficients
- 1024 channel 2 input dual polarization cross correlator (XMAC)
- 256 million channel PFB based spectrometer
- All optimizations were performed on the Simulink level, and tool flow options set from BEE\_ISE2, no HDL tweaking



# PFB1K (4 instances in 1 FPGA)

- Resource Utilization:
  - Flip Flops: 45,856 (69%)
  - LUTs: 14,816 (22%)
  - Slices: 25,380 (76%)
  - Block RAMs: 216 (65%)
  - MULT18X18s: 256 (78%)
- Max clock rate:
  - 252.8MHz (2VP70-7)
  - 72GMAC/s per FPGA @250MHz
- Power consumption: 26.5W
- Tool Flow run-time/Mem
  - Matlab/XSG: 10min/303MB
  - Synth: ~2 min/250MB
  - XFLOW: 84 min/1GB





# XMAC (5 instances in 1 FPGA)


- Resource Utilization:
  - Flip Flops: 45,765 (69%)
  - LUTs: 33,420 (50%)
  - Slices: 24,915 (75%)
  - Block RAMs: 0 (0%)
  - MULT18X18s: 0 (0%)
- Max clock rate:
  - 227.8MHz (2VP70-7)
  - 200 GMAC/s per FPGA @ 200MHz
- Power consumption: 6.65W
- Tool Flow run-time/Mem
  - Matlab/XSG: 2min/212MB
  - Synth: 2 min/178MB
  - XFLOW: 137 min/1.2GB





# 256M channel spectrometer

- 8K PFB (64K tap) + 32K pt FFT
- Resource Utilization:
  - Flip Flops: 30,964 (46%)
  - LUTs: 13,326 (20%)
  - Slices: 18,958 (57%)
  - Block RAMs: 328 (100%)
  - MULT18X18s: 128 (39%)
- Max Clock Rate:
  - 230.5 MHz (2VP70-7)
  - 28.8 GMAC/s per FPGA @200MHz
- Power Consumption: 13W
- Tool Flow run-time/Mem
  - Matlab/XSG: 32min/1.2GB
  - Synth: ~2 min/380MB
  - XFLOW: 64 min/890MB





# Xilinx FPGA vs. TI DSPs

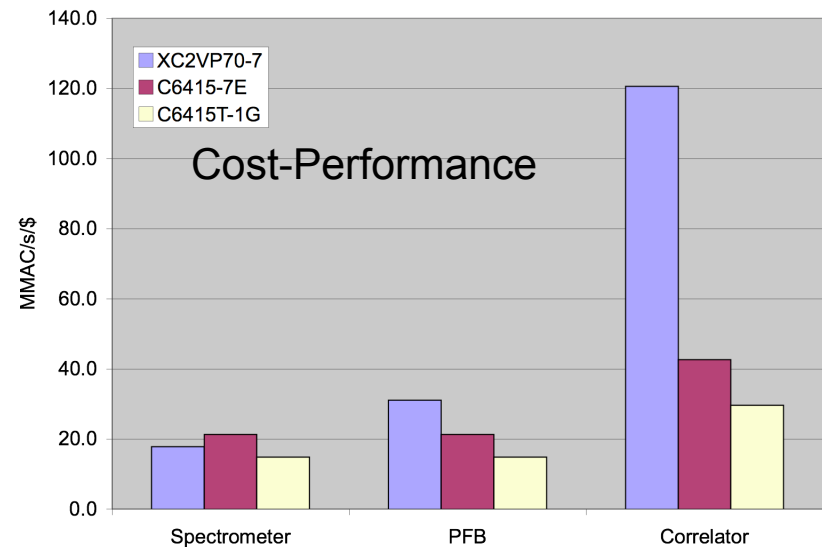

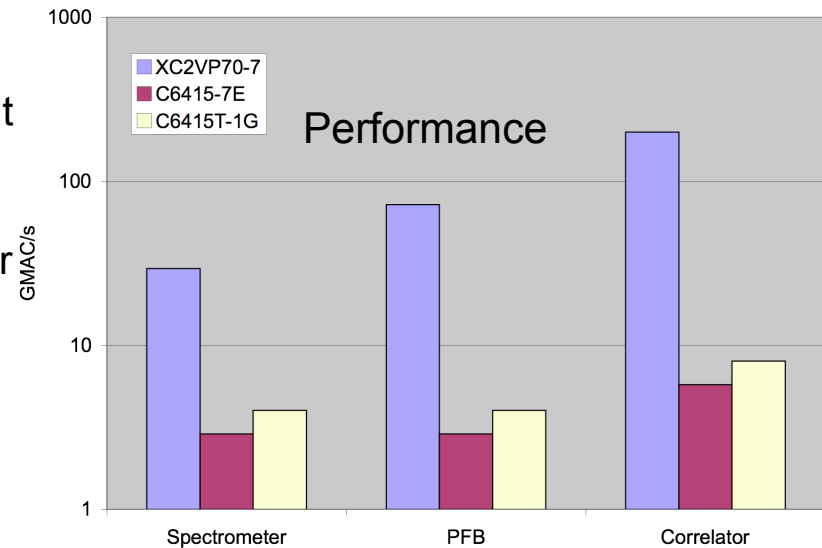
---

- Xilinx Virtex 2 Pro FPGA
  - 2VP70
    - Technology: 130nm CMOS
    - Clock rate: 200~250MHz
    - Unit Cost: \$1500~2400
  - 2VP20
    - Technology: 130nm CMOS
    - Clock rate: 200~250MHz
    - Unit Cost: \$366
- TI C640 DSP
  - C6415T-1G
    - Technology: 90nm CMOS
    - Clock rate: 1GHz
    - Unit Cost: \$270
  - C6415-7E
    - Technology: 130nm CMOS
    - Clock rate: 720MHz
    - Unit Cost: \$135




# Comparison with DSP Chips

- Spectrometer & polyphase filter bank (PFB): 16 mult, 32bit acc, Correllator: 4bit mult, 24bit acc.
- Cost based on street price.
- Assume peak numbers for DSPs, mapped for FPGAs.
- TI DSPs:
  - C6415-7E, 130nm (720MHz)
  - C6415T-1G, 90nm (1GHz)
- FPGAs
  - 130nm, freq. 200-250MHz.
- **Metrics include chips only (not system)**





# Cost Effective Solutions for PFB





# Project status

---

- 10 node system manufacturing (8/2005)
- Demonstration applications:
  - NASA DSN 128M channel spectrometer (7/2005)
  - VLBI 1GHz spectrum data recorder (9/2005)
  - 8 antenna 200MHz dual polarization correlator (9/2005)