The Design and Applications of BEE2: A High End Reconfigurable Computing System

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

Outline

- Motivations for High End Reconfigurable Computing
- BEE2 system
  - Hardware architecture
  - Programming environment
- Demonstration applications
  - SETI billion channel spectrometer
  - Antenna array correlator
- Current status

High-End Reconfigurable Computer (HERC)

- A computer with *supercomputer-like performance*, based solely on FPGAs and/or other reconfigurable devices as the processing elements.
- 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)

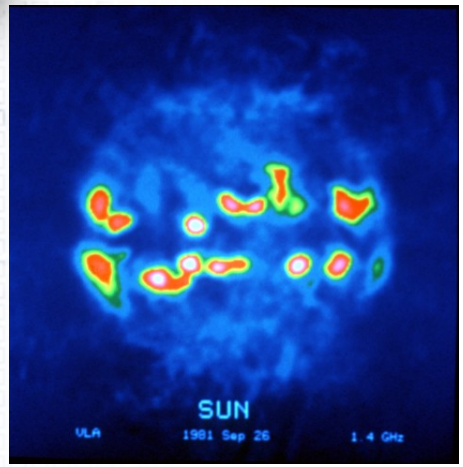

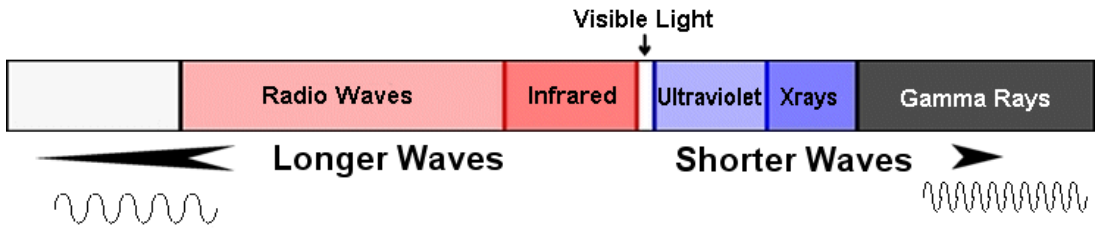

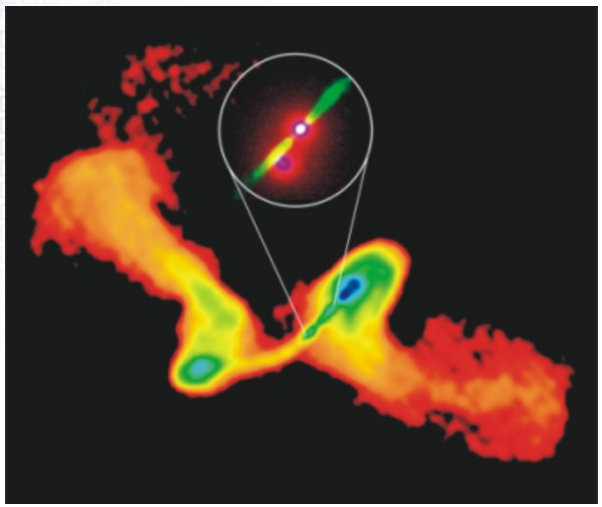

Colliding black holes

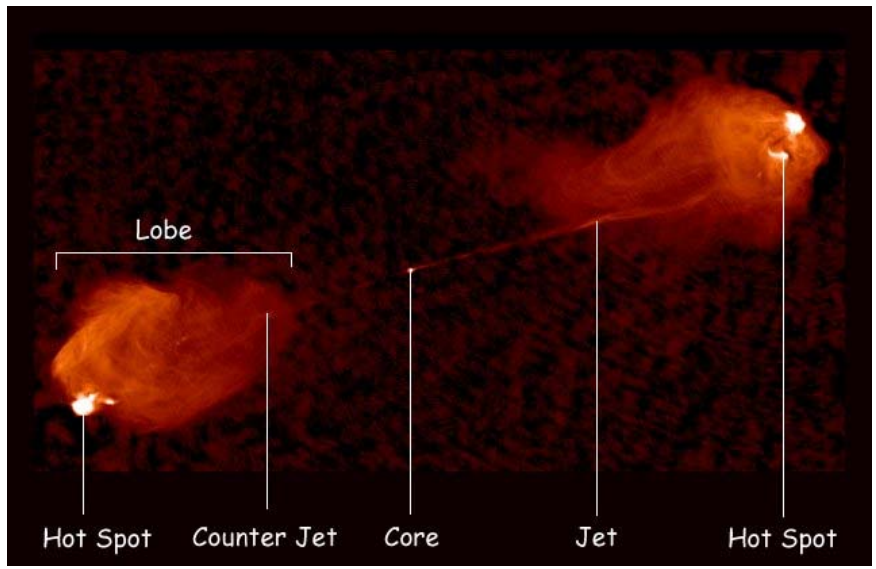

Radio jets

Image courtesy of NRAO/AUI

Allen Telescope Array (CA, USA)

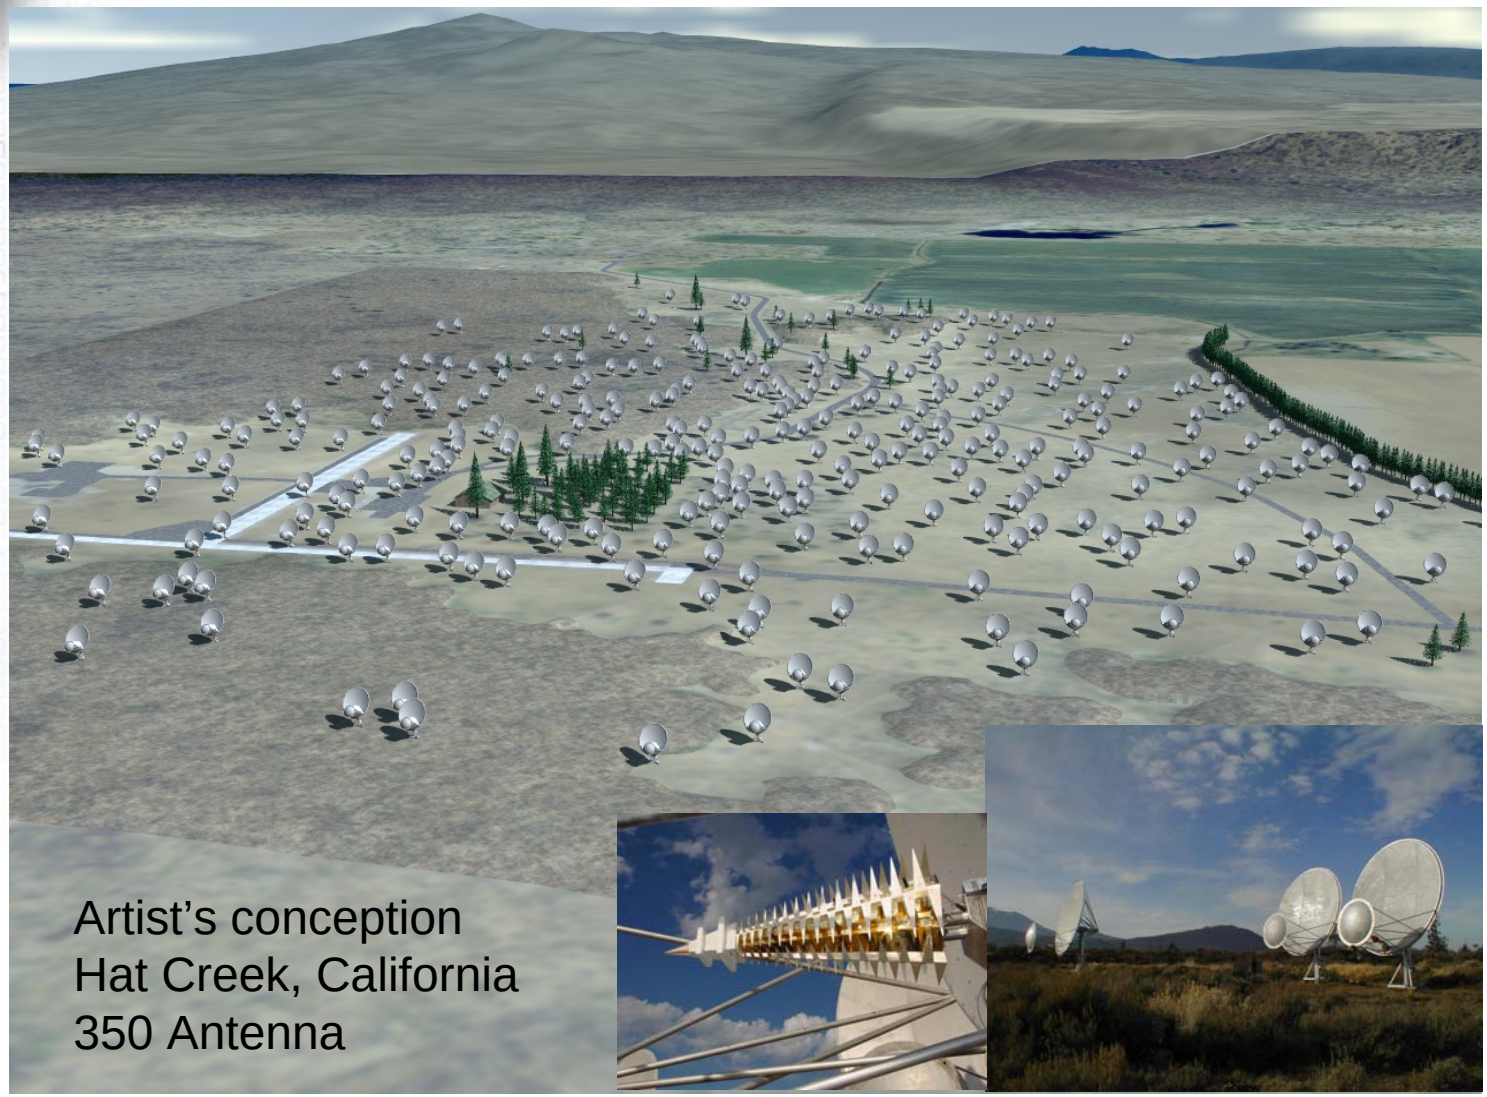

Large-N, Small-D Concept

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

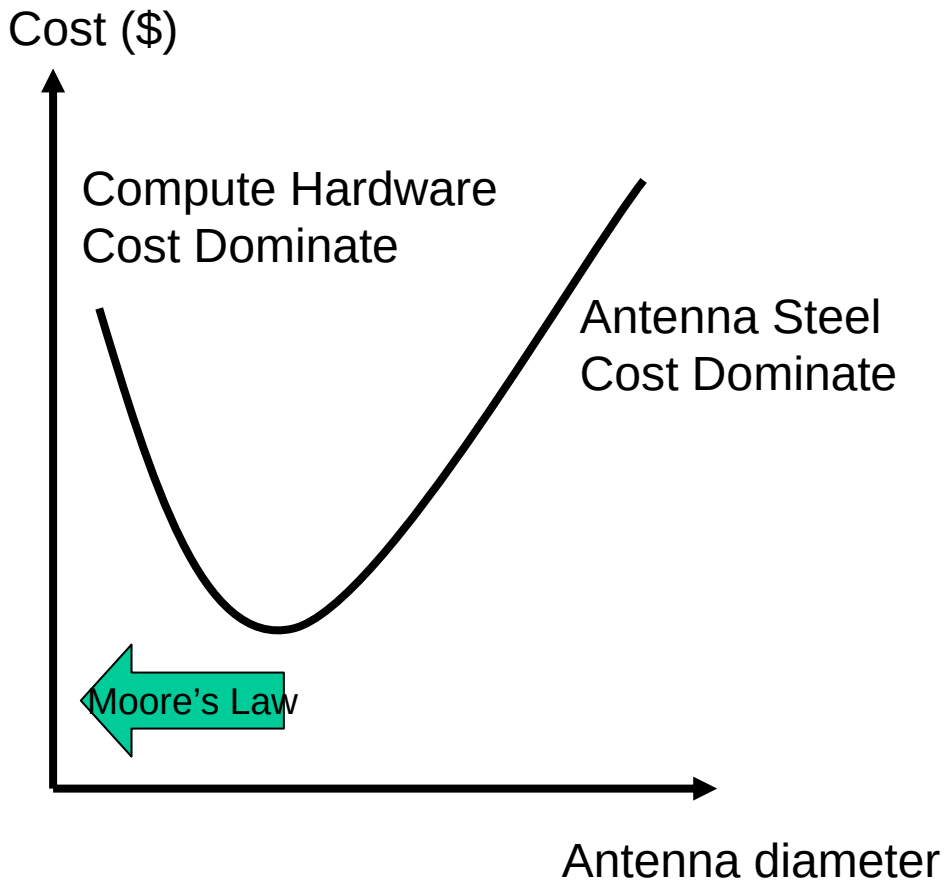

Problems with existing approach

- 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

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

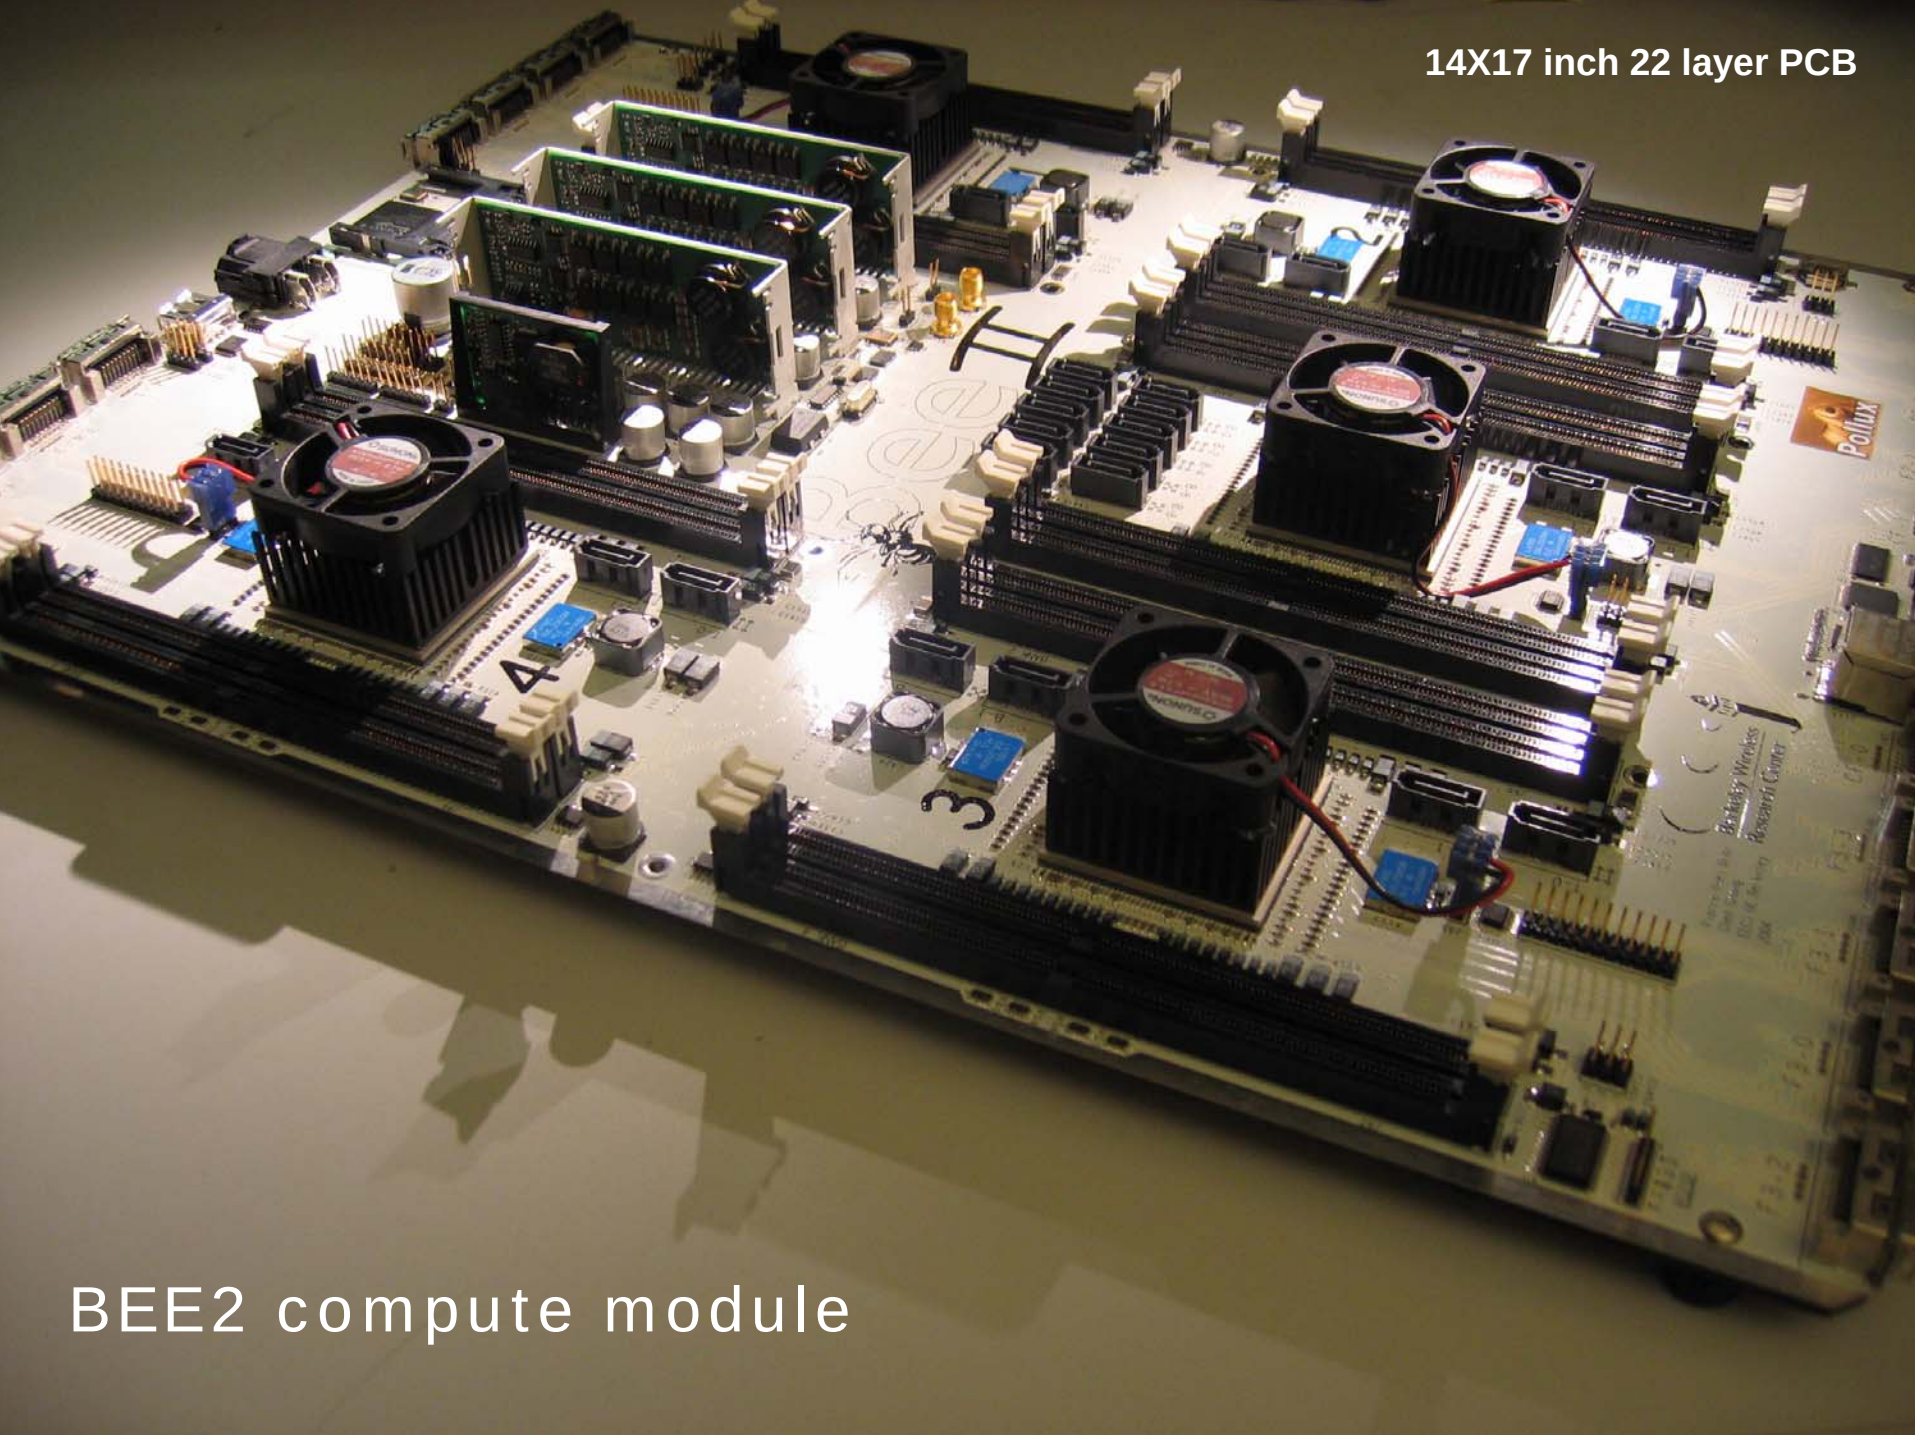
14X17 inch 22 layer PCB
BEE2 compute module

Compute Module Diagram

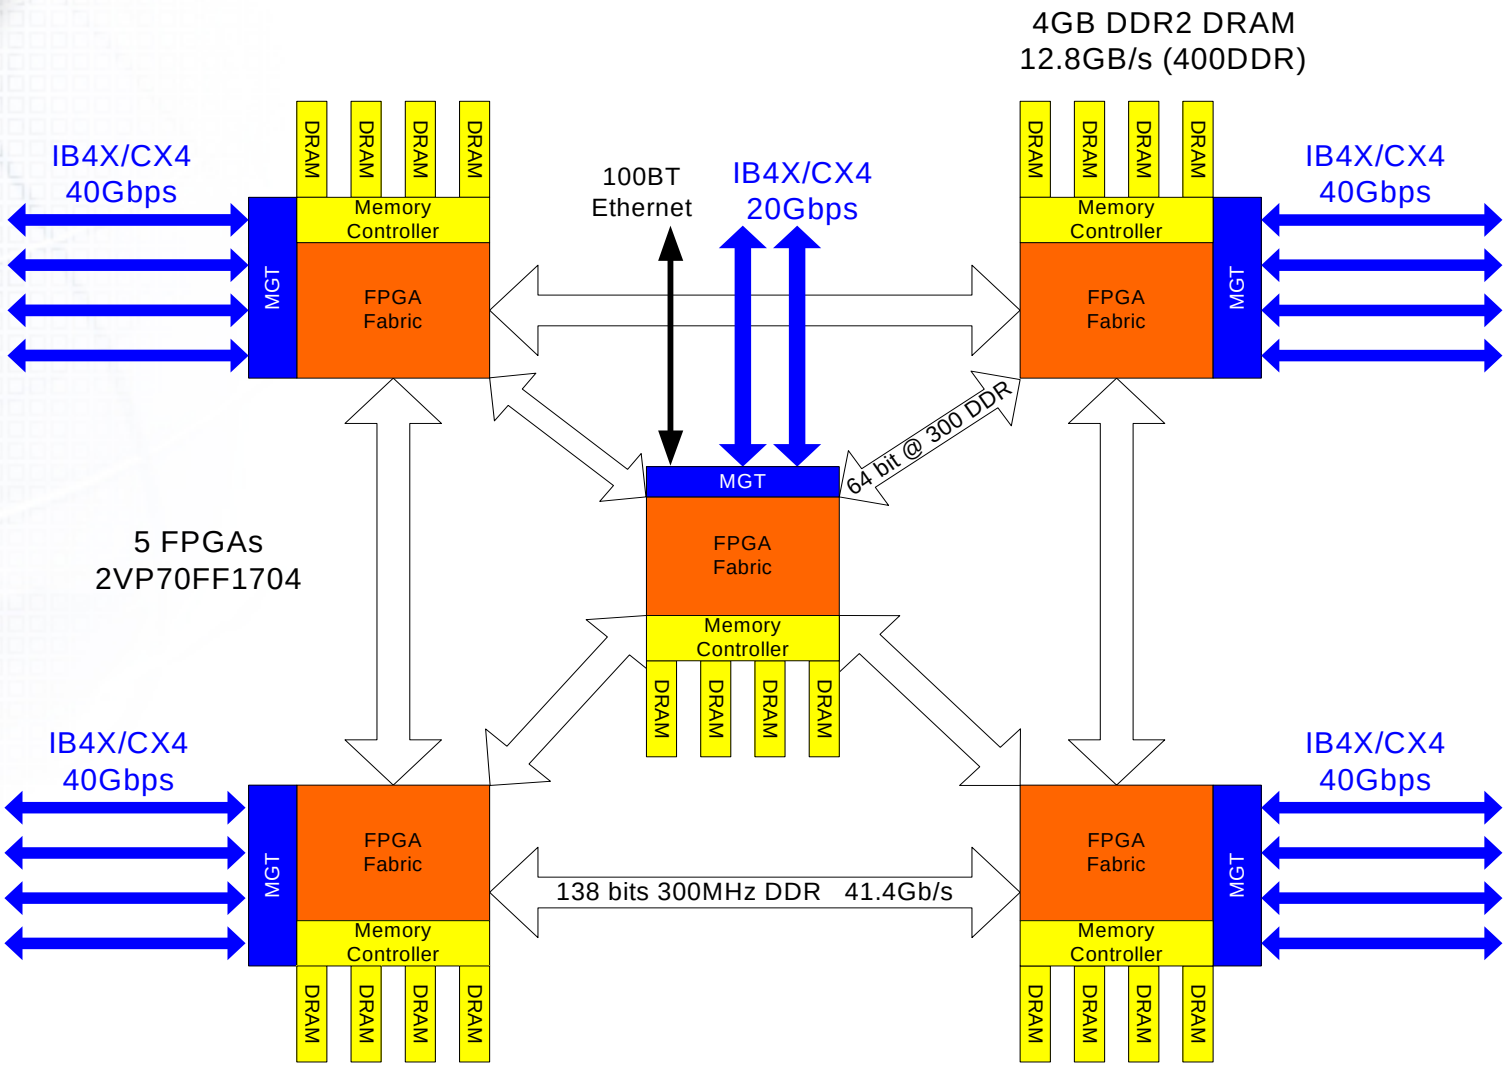

Inter-Module Connections

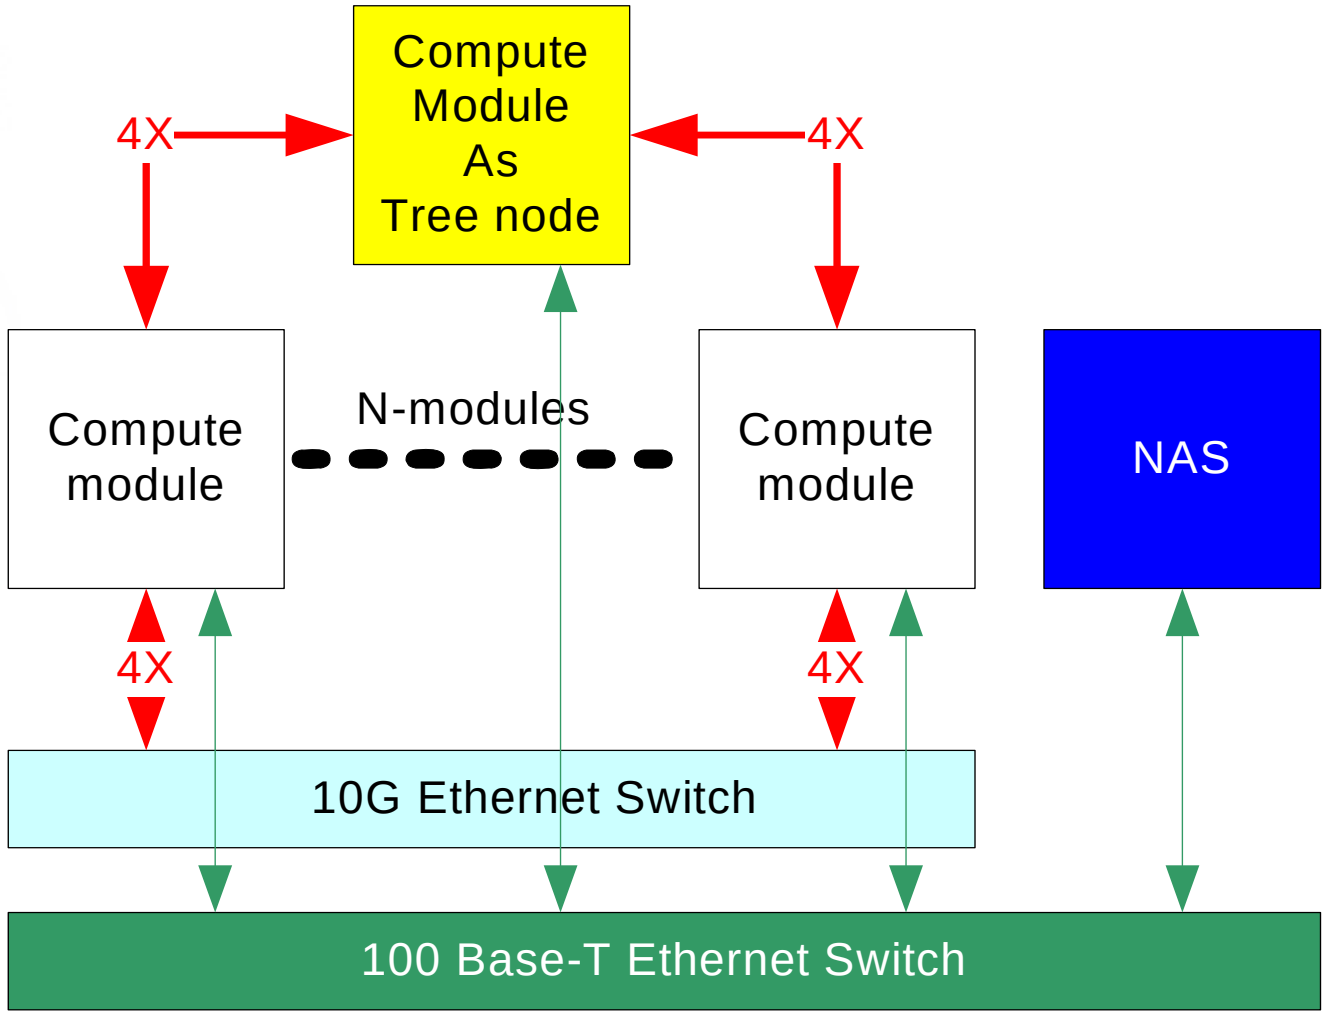

19” Rack Cabin Capacity

- 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

BEE2 A/D interface overview

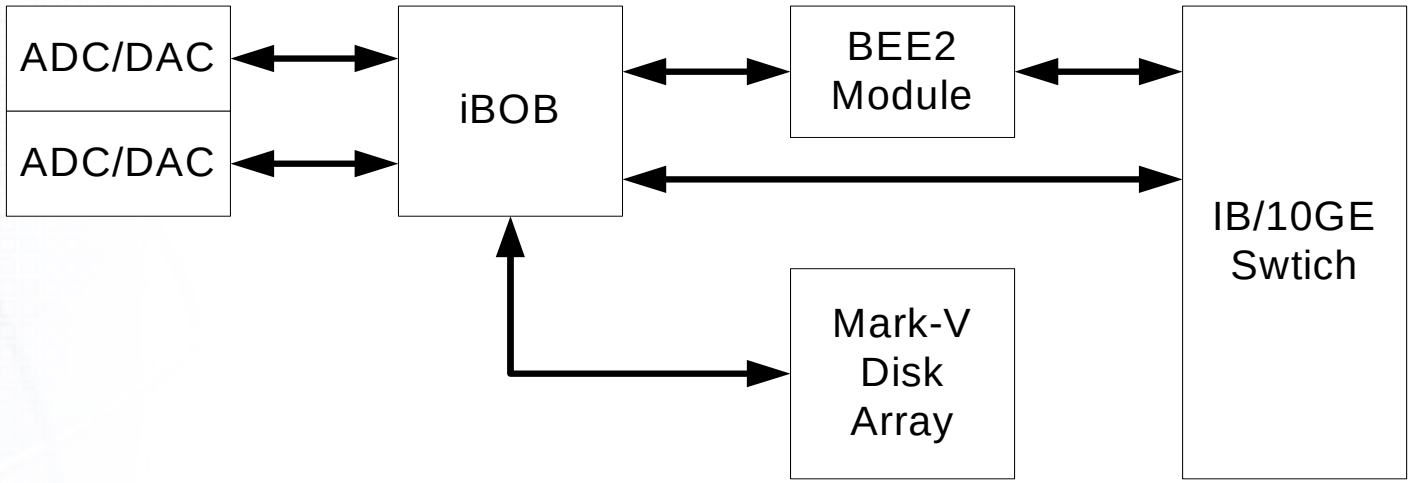

- Use iBOB to fanout the 10G IB4X serial connections to parallel LVDS/LVPEL signals
- iBOB can be connected to BEE2 modules or directly to Infiniband or 10GE switches
- Built-in support to connect to the Mark-V disk array archiver

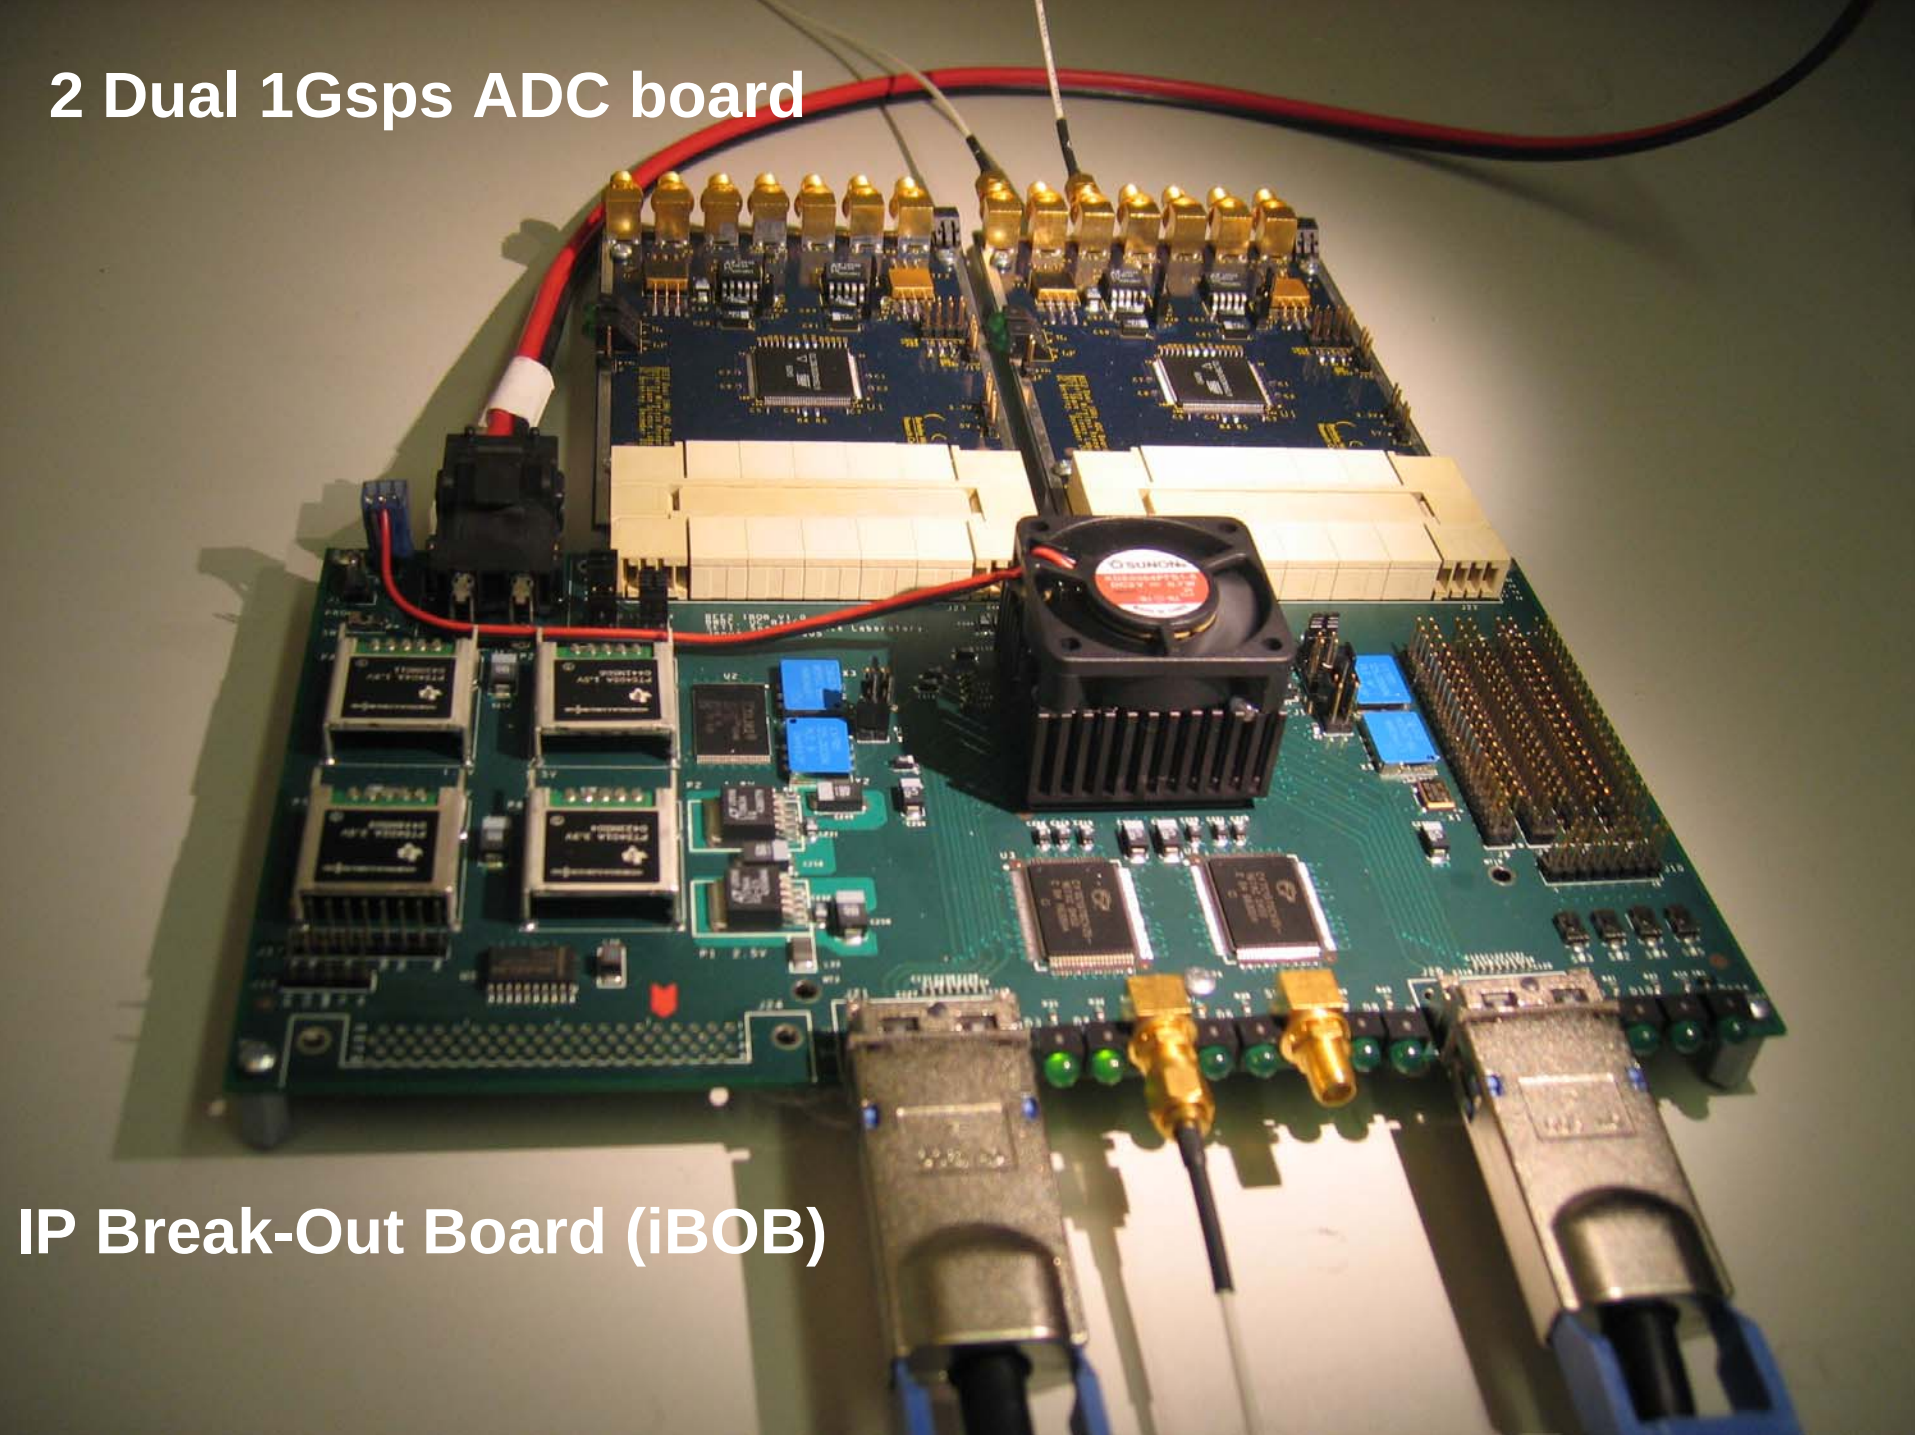
2 Dual 1Gsps ADC board
IP Break-Out Board (iBOB)

BEE2 Design Environment

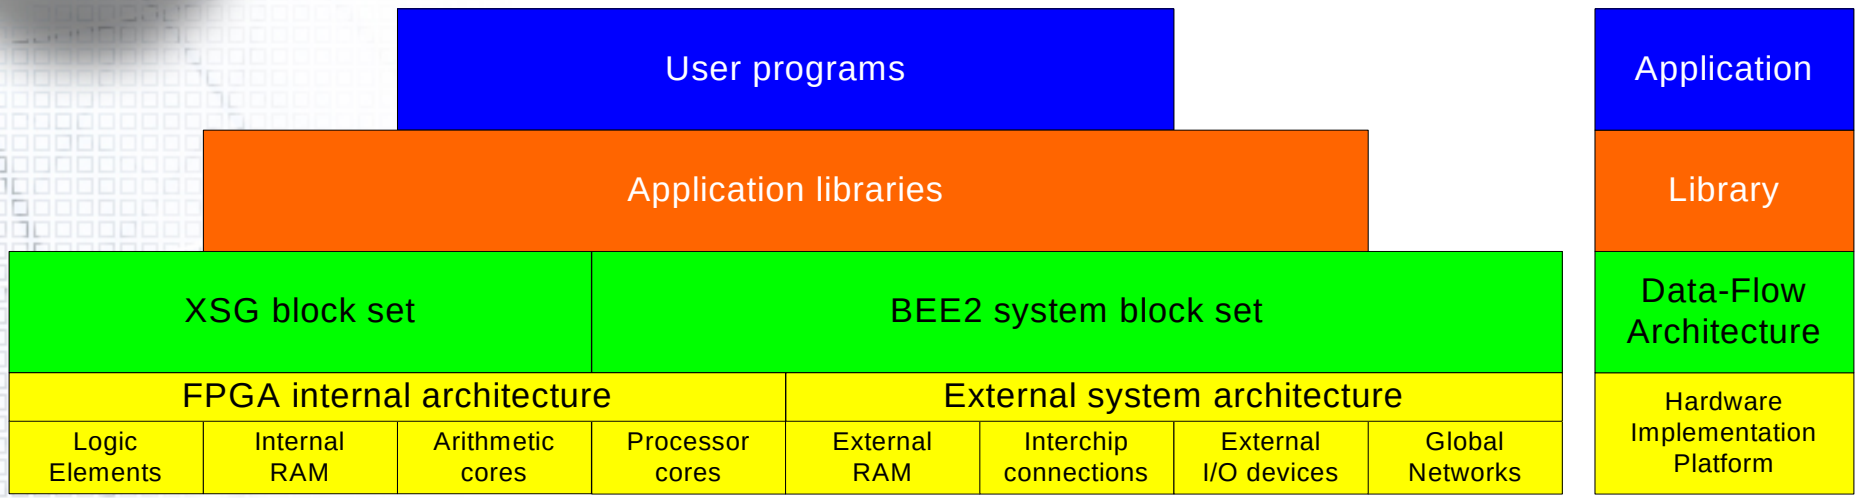

- Layered design abstractions
- Platform-based hardware abstraction
- Discrete-time data flow execution model
- Script-based dynamic application library generation
- Hardware/software codesign with full system integration

BEE Platform Studio (BPS)

- BEE Platform Studio 1.1
  - Simulink library support for system devices
  - Full SoC generation include software integration
  - Simple GUI in Matlab
  - Require no knowledge of backend tool flow from end users
- Backend commercial tools
  - Xilinx ISE 7.1i SP4
  - Xilinx EDK 7.1i sp2
  - Xilinx System Generator 7.1i
  - Synplicity Synplify Pro 8.1

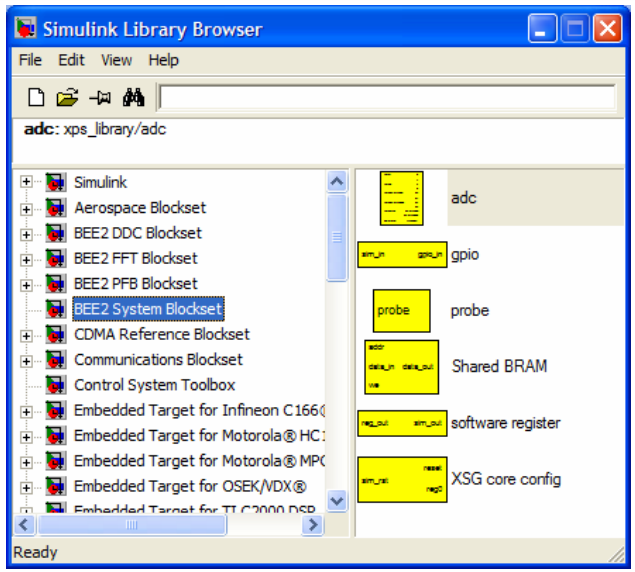

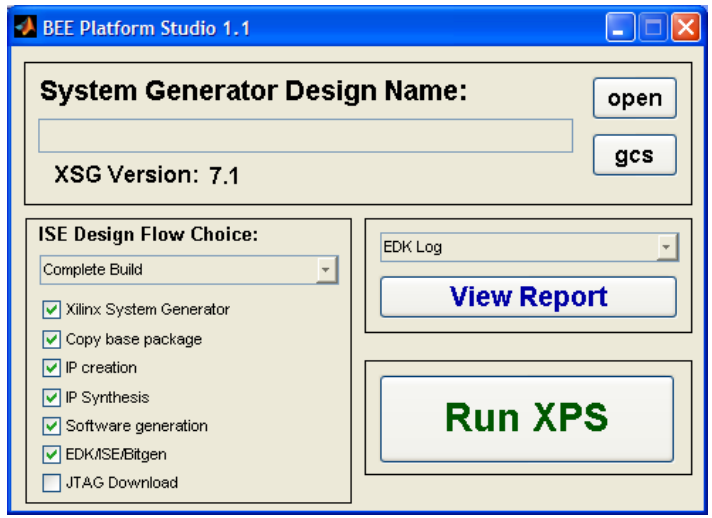

BPS design example

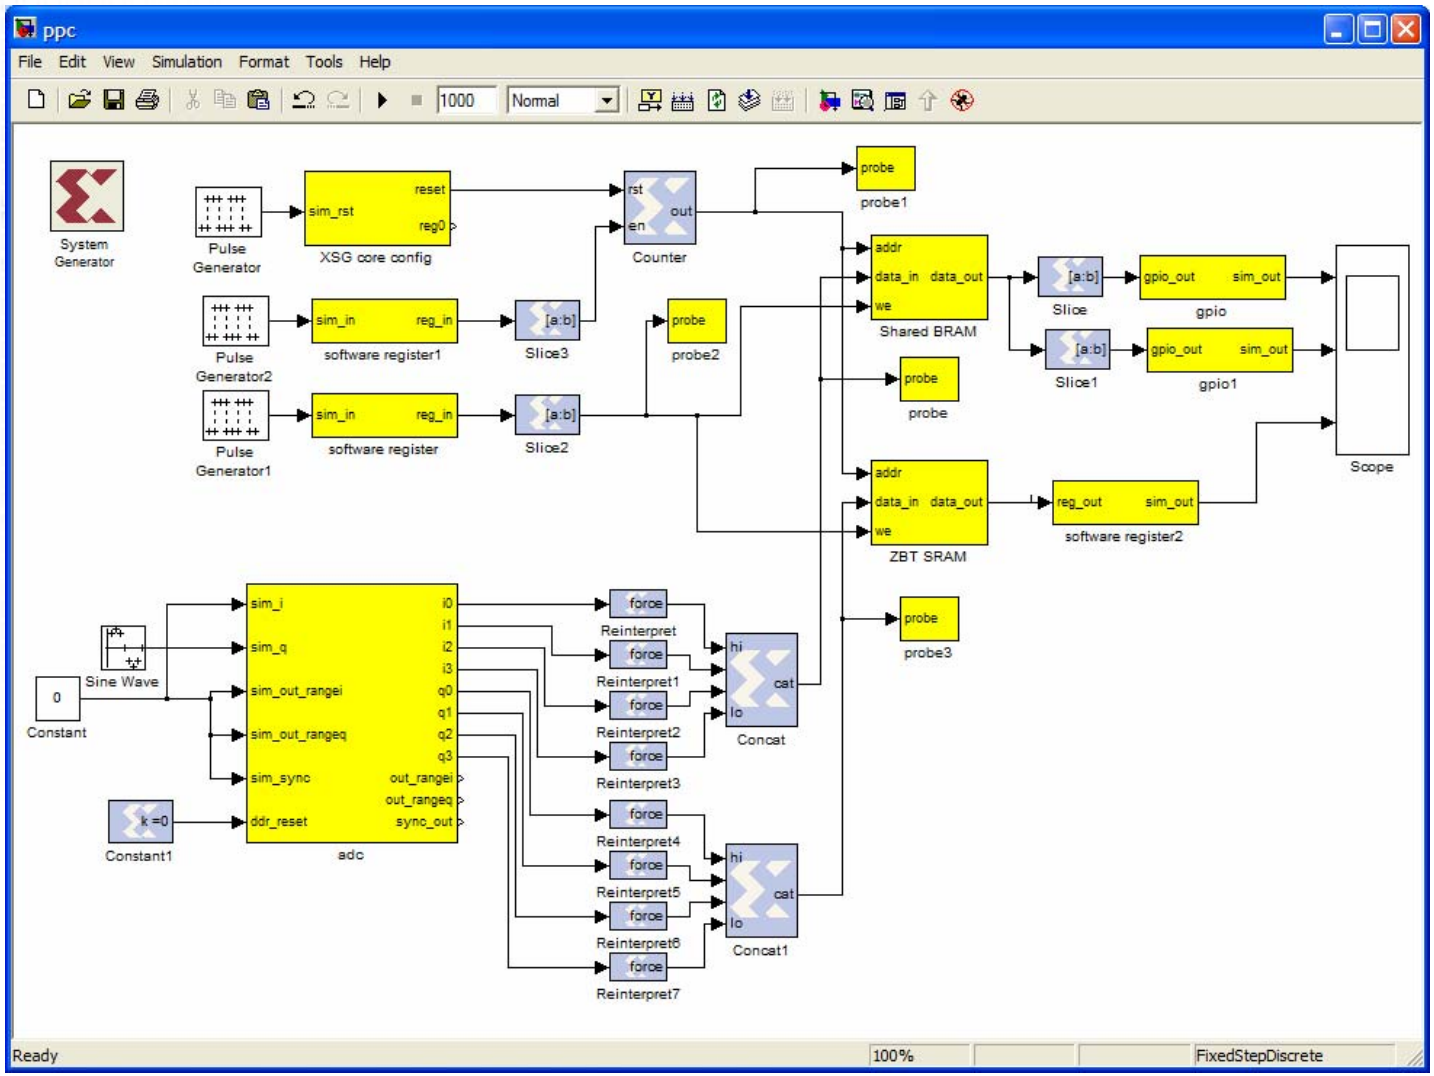

SETI Billion Channel Spectrometer

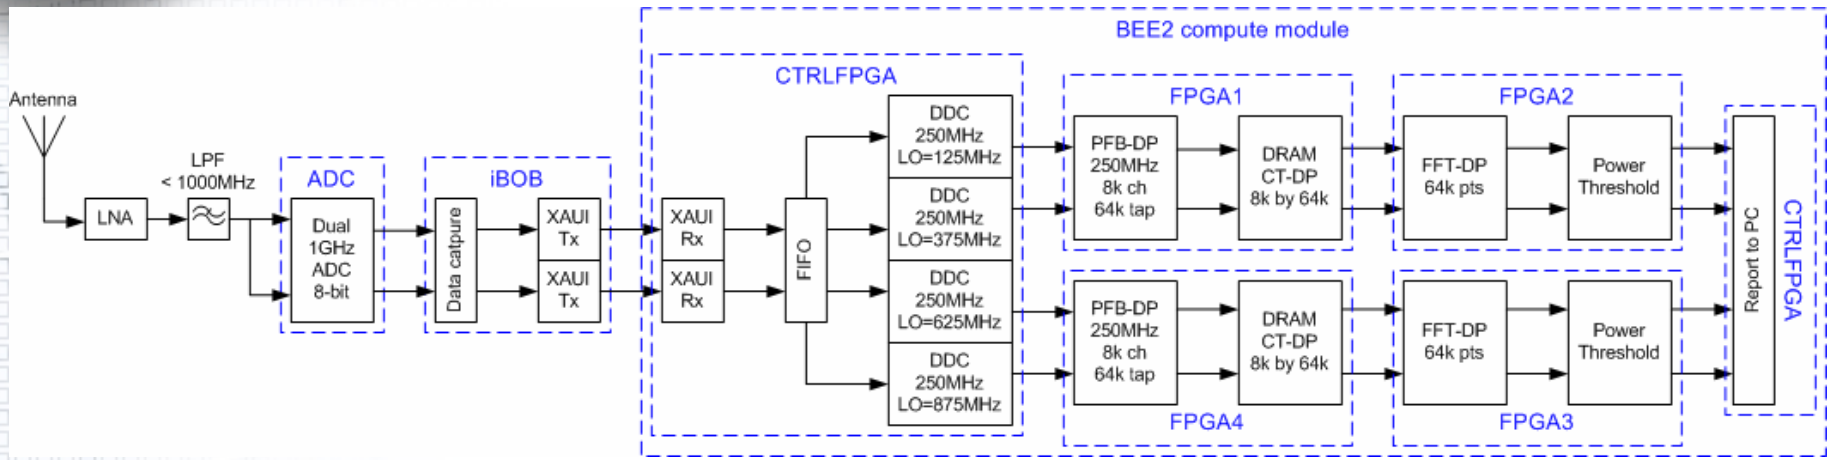

- 0.7Hz channels over 800MHz ➔ 1 billion Channel real-time spectrometer
- Implemented on 1 BEE2 module and yields 333GOPS (16-bit mults, 32-bit adds), at 150Watts (similar to desk-top computer)
- >100x peak throughput of current Pentium-4 system on integer performance, & >100x better throughput per energy.

JPL/SETI Galactic Plane Sky Survey

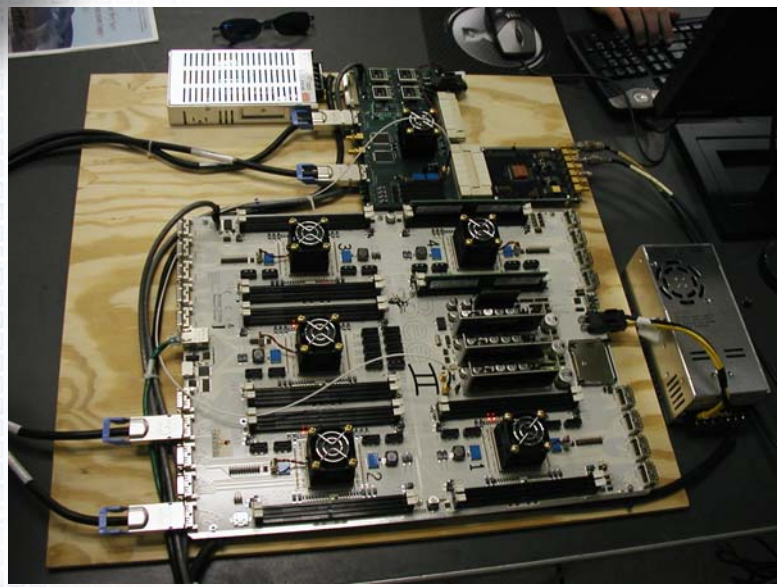

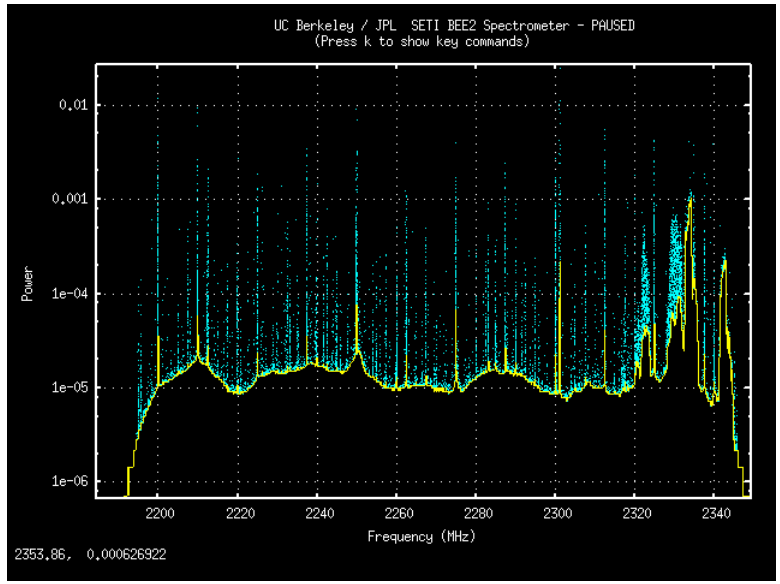

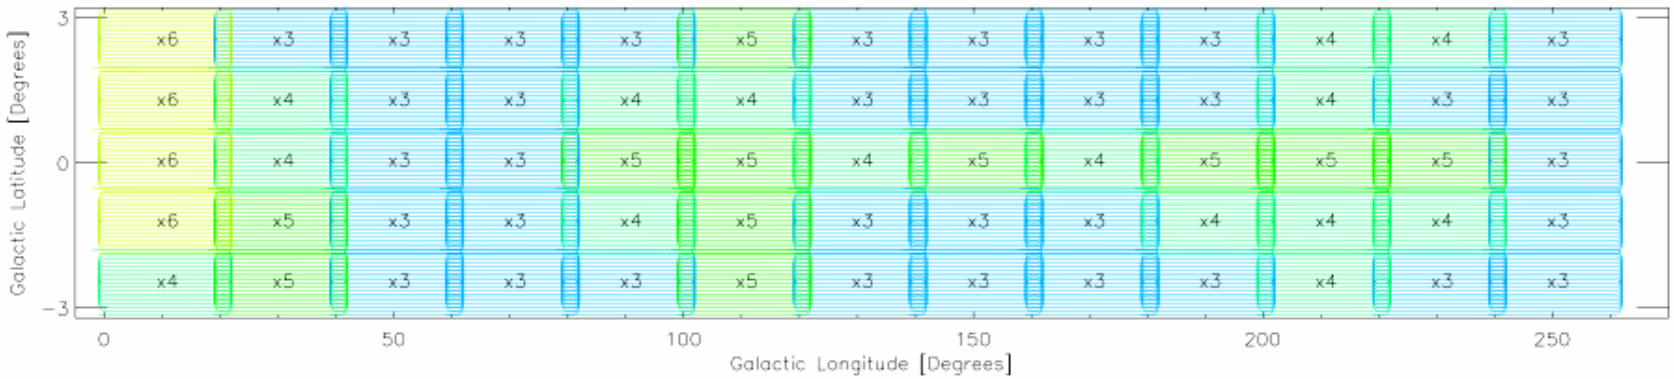

Antenna array correlation basics

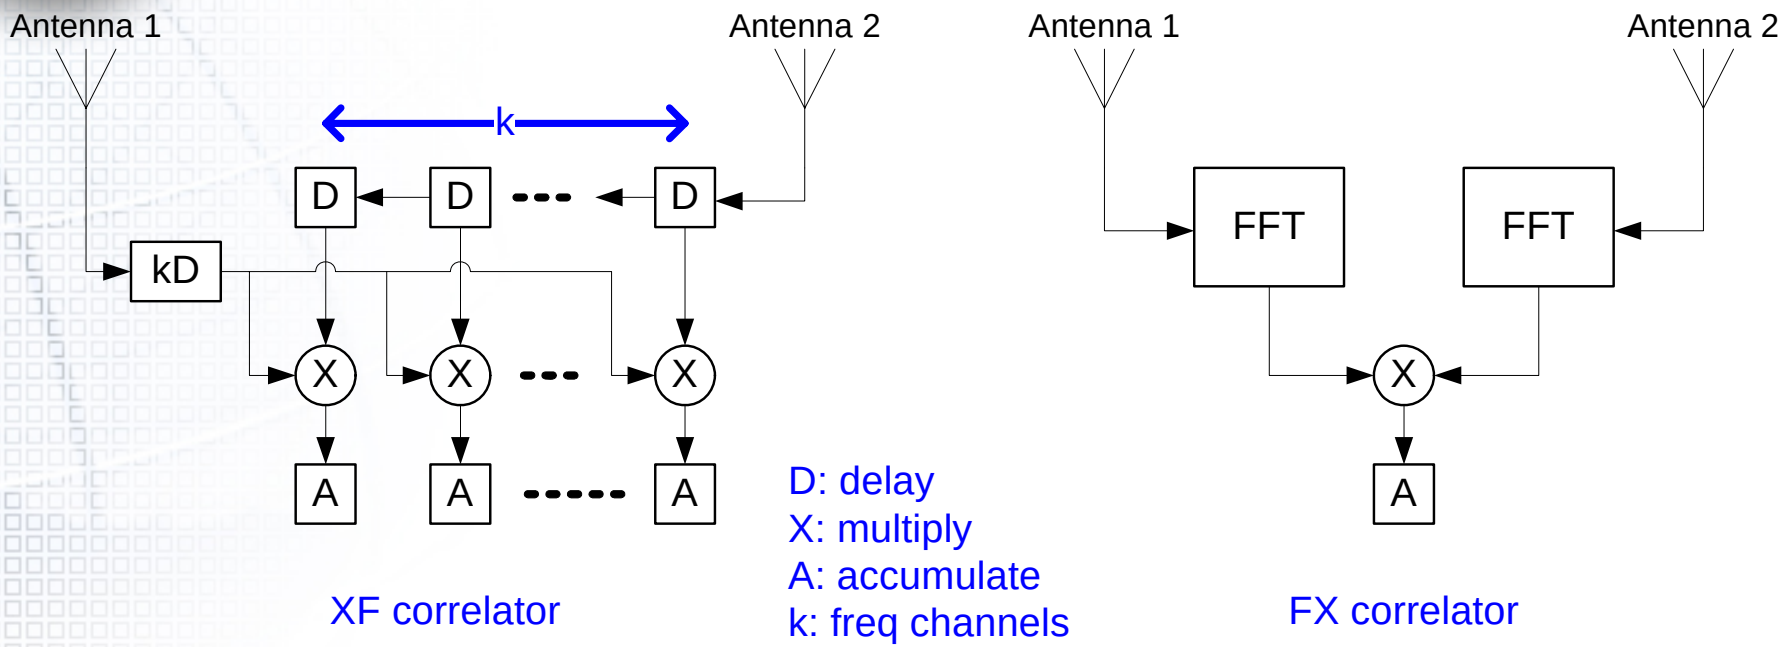

- FX correlator is orders of magnitude more computationally efficient

N-antenna correlator overview

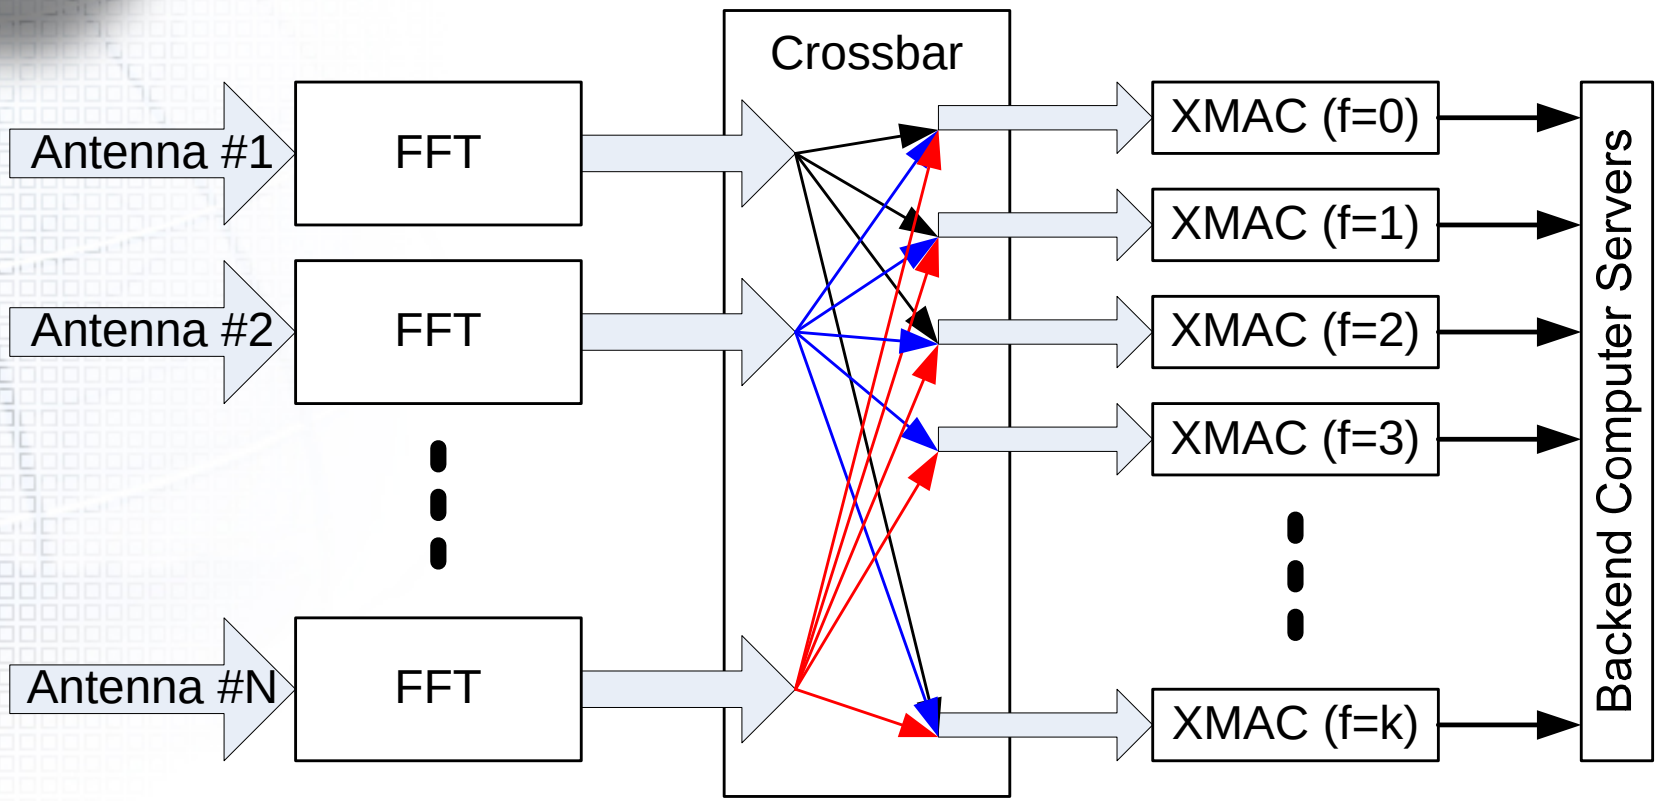

For 350 antenna, each 100MHz bandwidth require over 100 TOPS compute throughput

Frequency transform

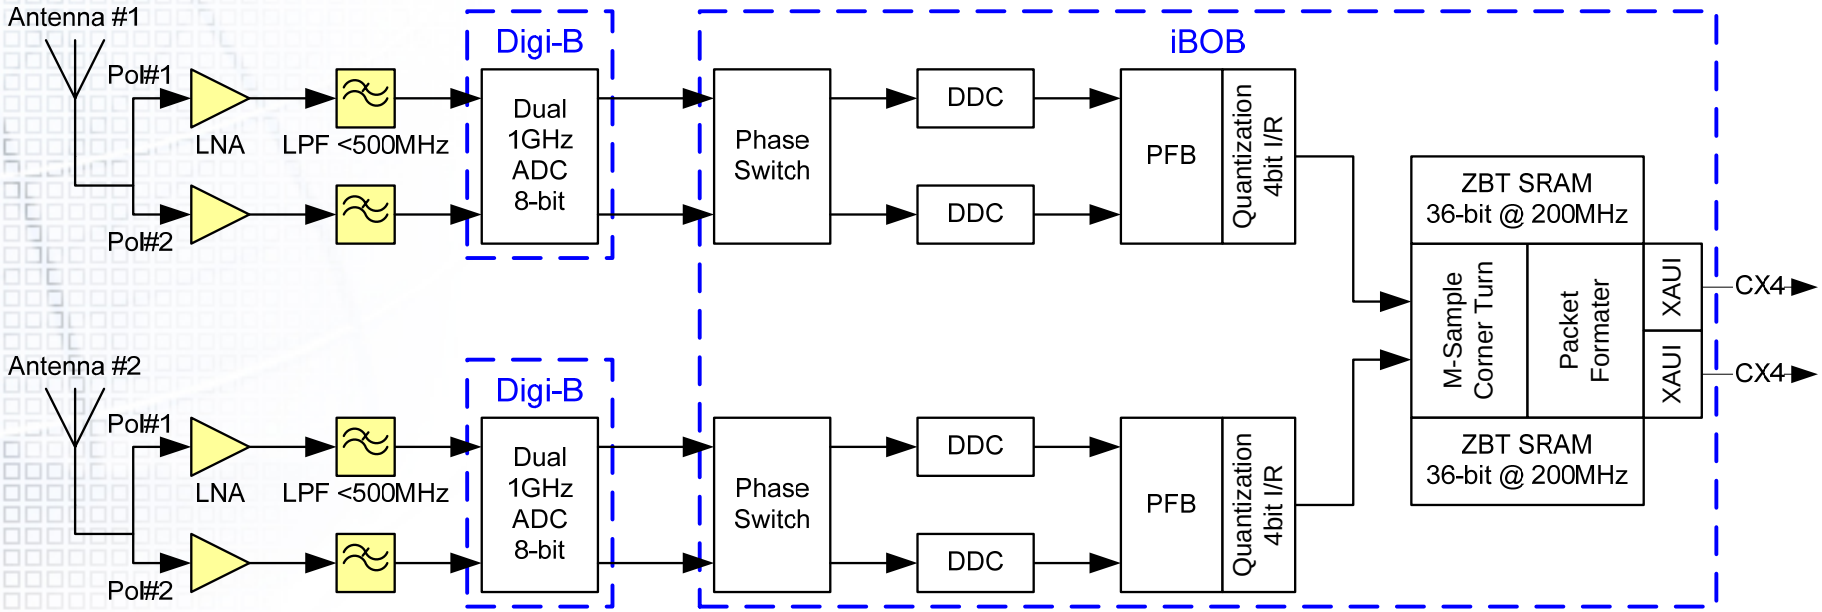

XMAC implementation

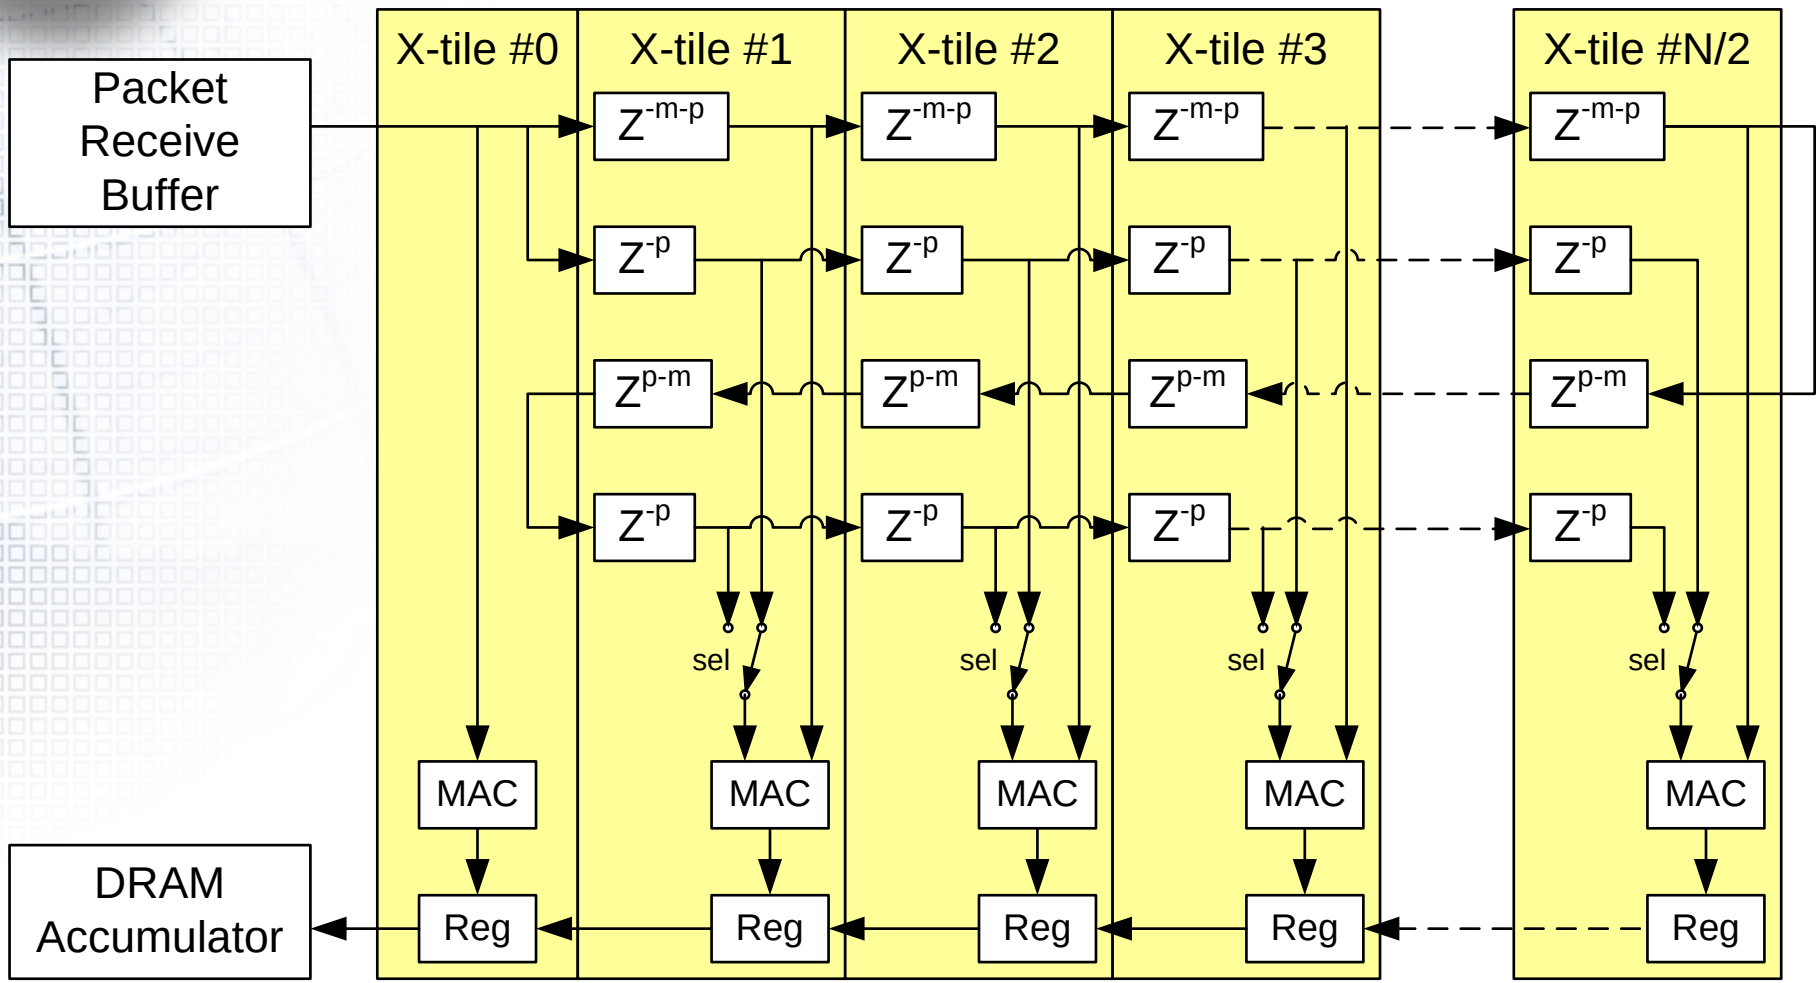

Crossbar switch implementation

- For small number of antennas (N<16)
  - Frame based circuit switching using BEE2 module internal resources
- For large number of antennas (N>16)
  - Packet switching using commercial 10G Ethernet switches

Frequency Transform Outputs				Switch	XMAC inputs			
A0F1	A0F2	A0F3	A0F0		A3F0	A2F0	A1F0	A0F0
A1F2	A1F3	A1F0	A1F1		A0F1	A3F1	A2F1	A1F1
A2F3	A2F0	A2F1	A2F2		A1F2	A0F2	A3F2	A2F2
A3F0	A3F1	A3F2	A3F3		A2F3	A1F3	A0F3	A3F3
T=3	T=2	T=1	T=0		T=3	T=2	T=1	T=0

200MHz 4 antenna correlator on a single iBOB

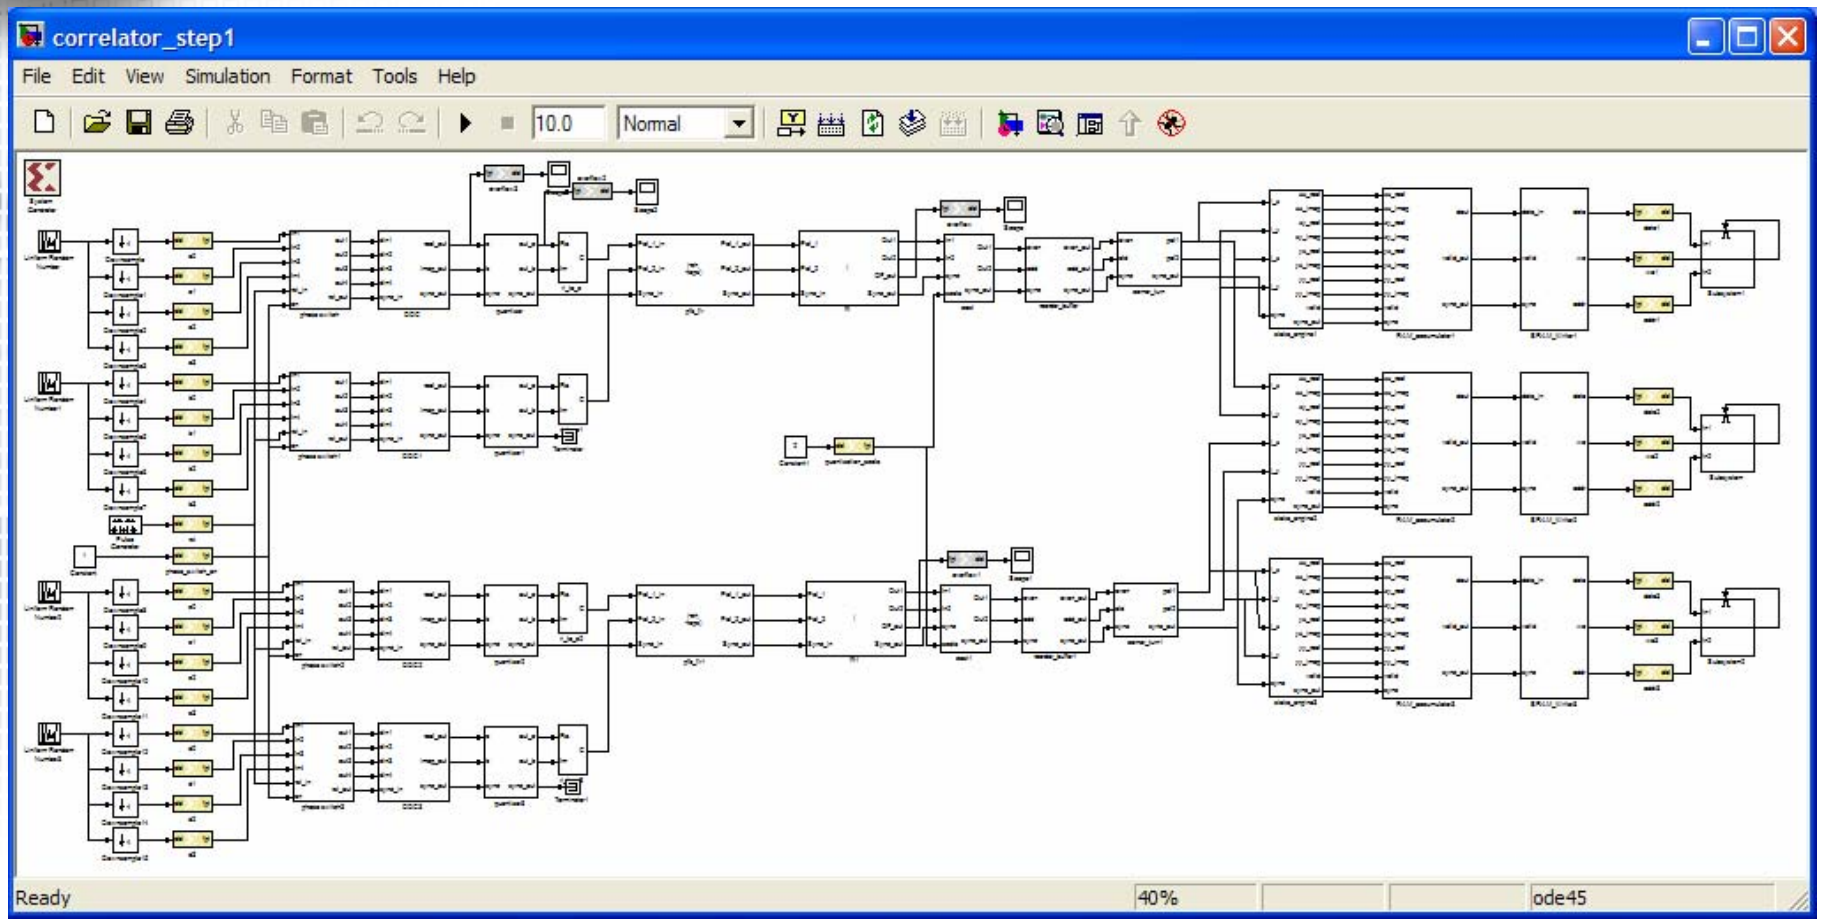

Field test at CARMA

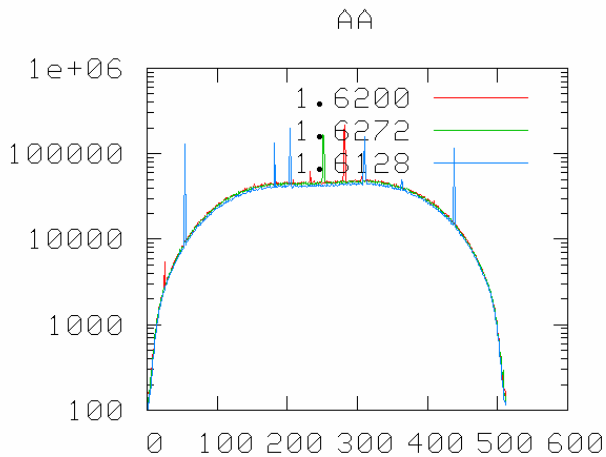

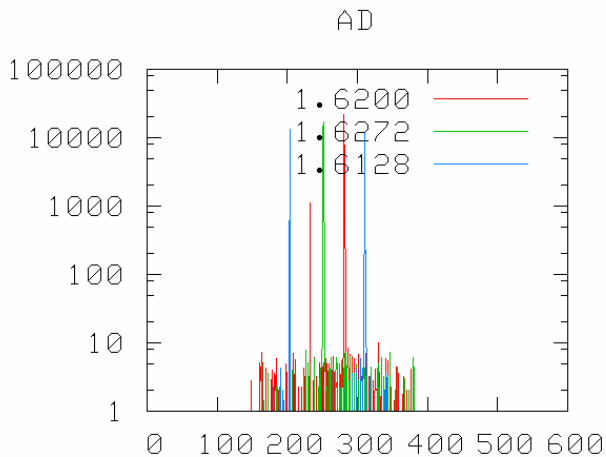

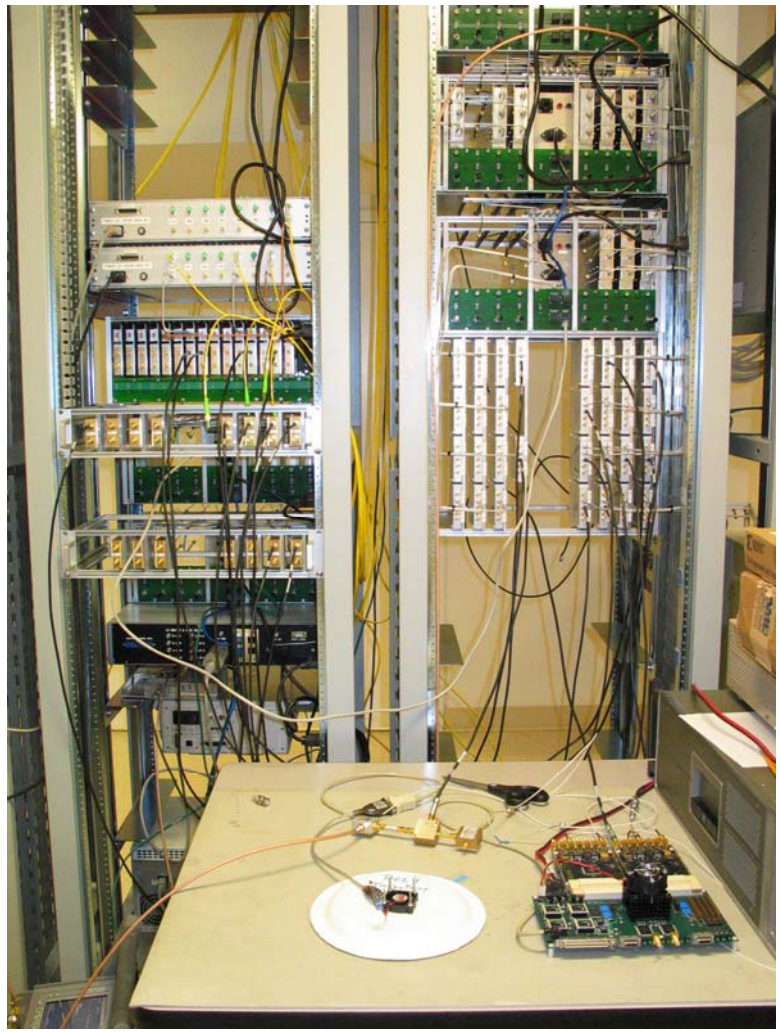

Project status

- 10 node system manufacturing (10/2005)
- Demonstration applications:
  - NASA DSN 128M channel spectrometer (8/2005)
  - 1GHz portable real-time spectrometer (10/2005)
  - VLBI 1GHz spectrum data recorder (12/2005)
  - 8 antenna 200MHz dual polarization correlator (12/2005)