UT-BATTELLE

Presented at the OASCR Applied Mathematics PI Meeting
(Panel on Next Generation Architectures)

Steve Poole
Chief Scientist / Director of Special Programs
Computer Science and Mathematics Division

May 23, 2007
Livermore, CA.

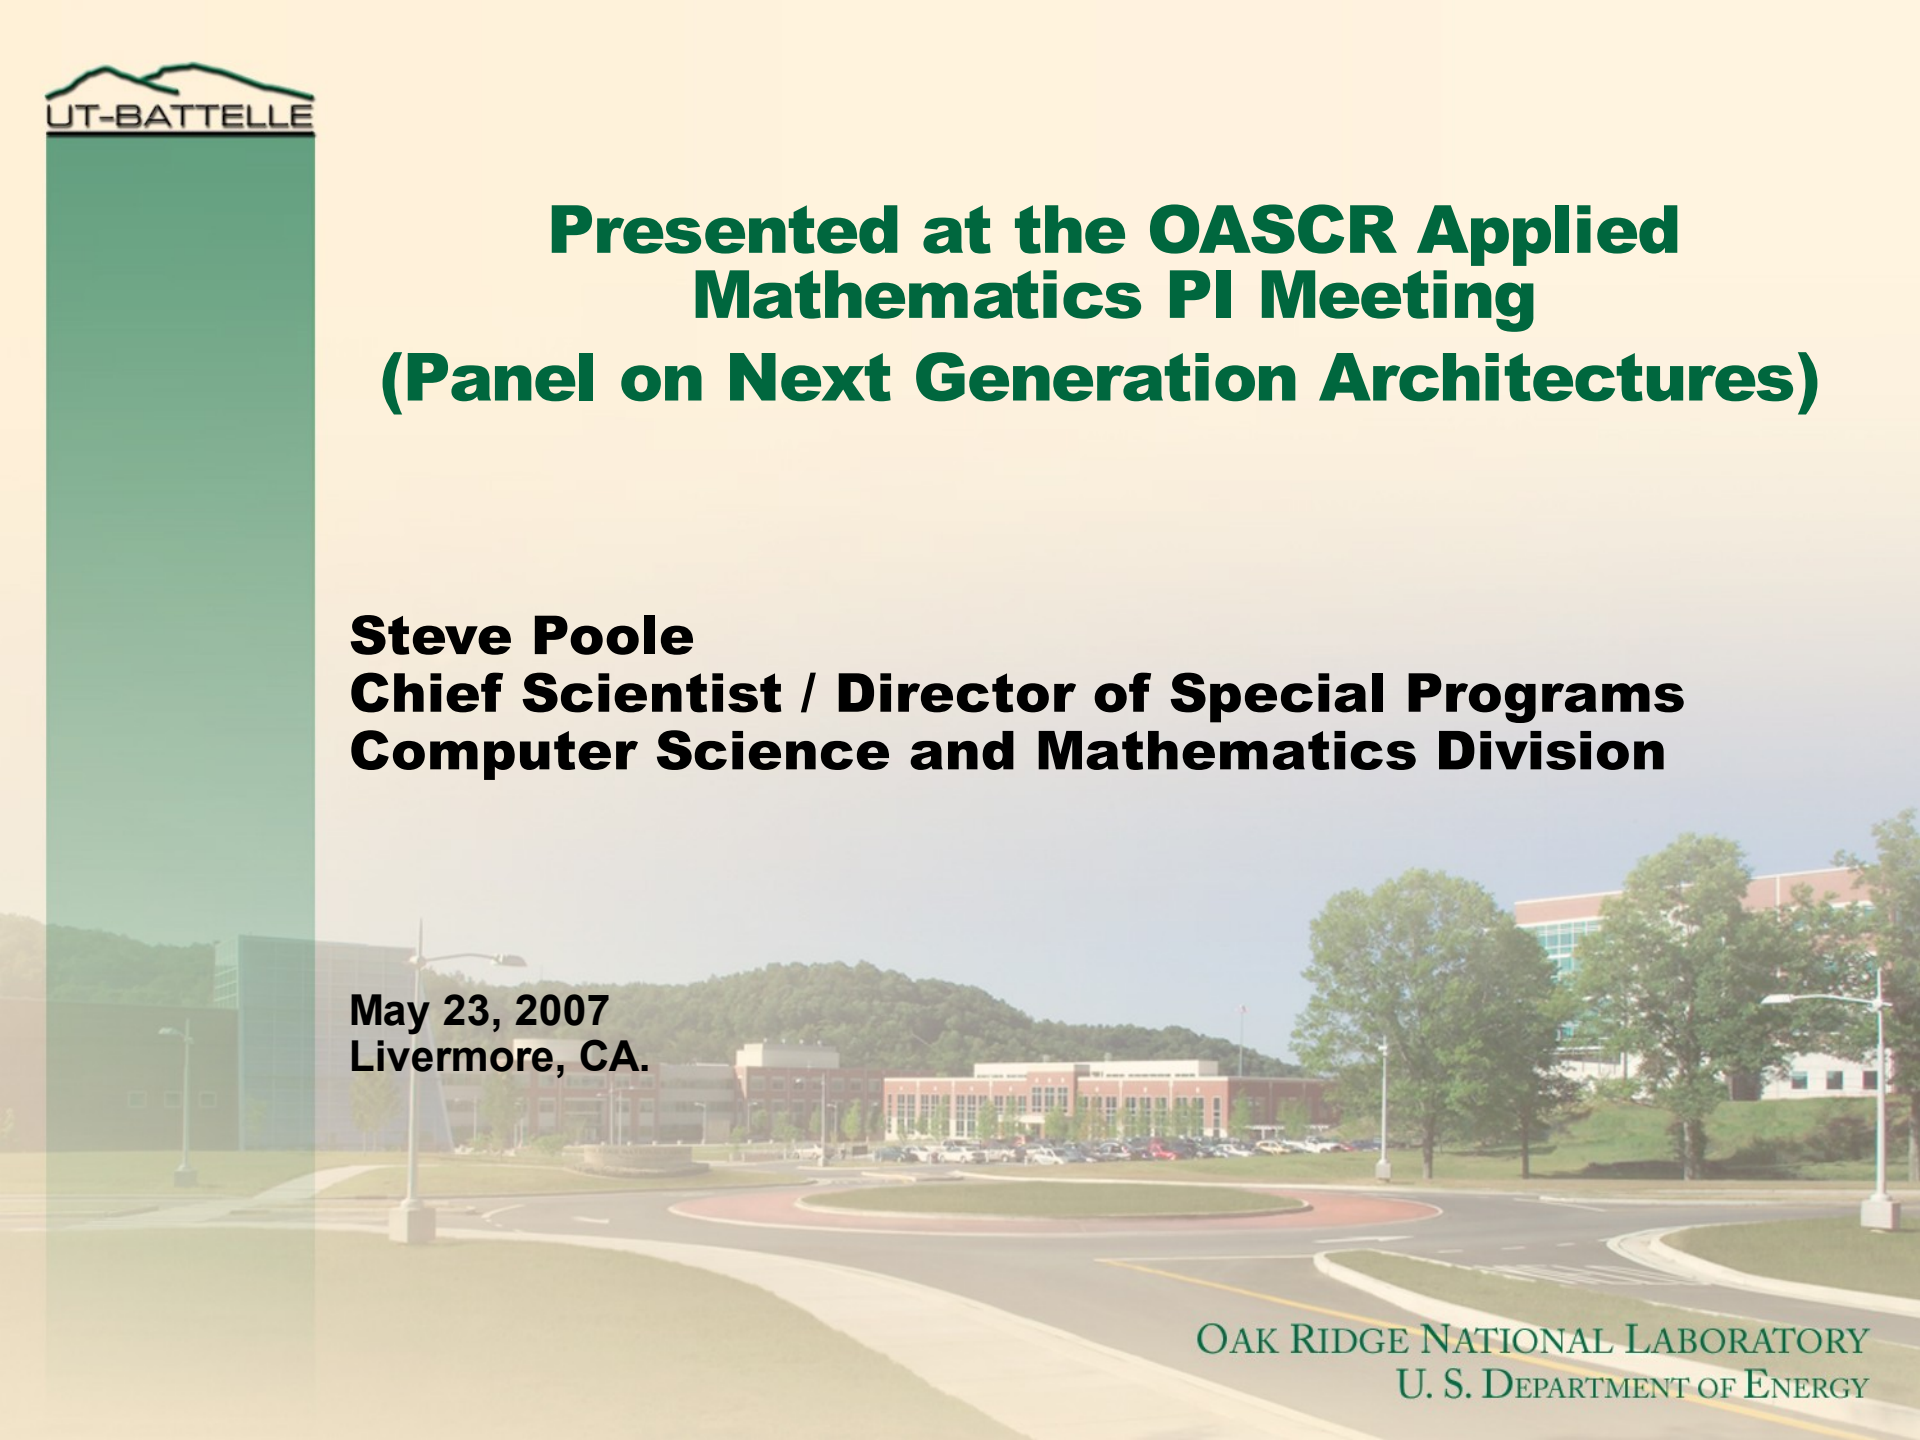

From yesterday & the Townhall meetings

What goes around, comes around (technology cycles)

Difficult in technologies to determine what is a Disruptive Technology
- Time based problem, 1997/1998, now
- Difficult to get funding

Rewriting/redesigning complete applications from first principles is very expensive
- But improved algorithms can make a HUGE difference

Contrary to popular beliefs, without software, there is not much use for computing hardware.

FLOPs are easy (useful flops are not)

We can no longer ride the coat tails of ASC(I) (HW Funding)

DOE-SC will have to learn about NRE (technologies)

We were asked to design an ExaFLOP system for 2015

http://computing.ornl.gov/workshops/town_hall/index.shtml

http://hpcrd.lbl.gov/E3SGS/main.html

Potential Systems Roadmap (2$$8-2$14)
Where are your algorithms ?

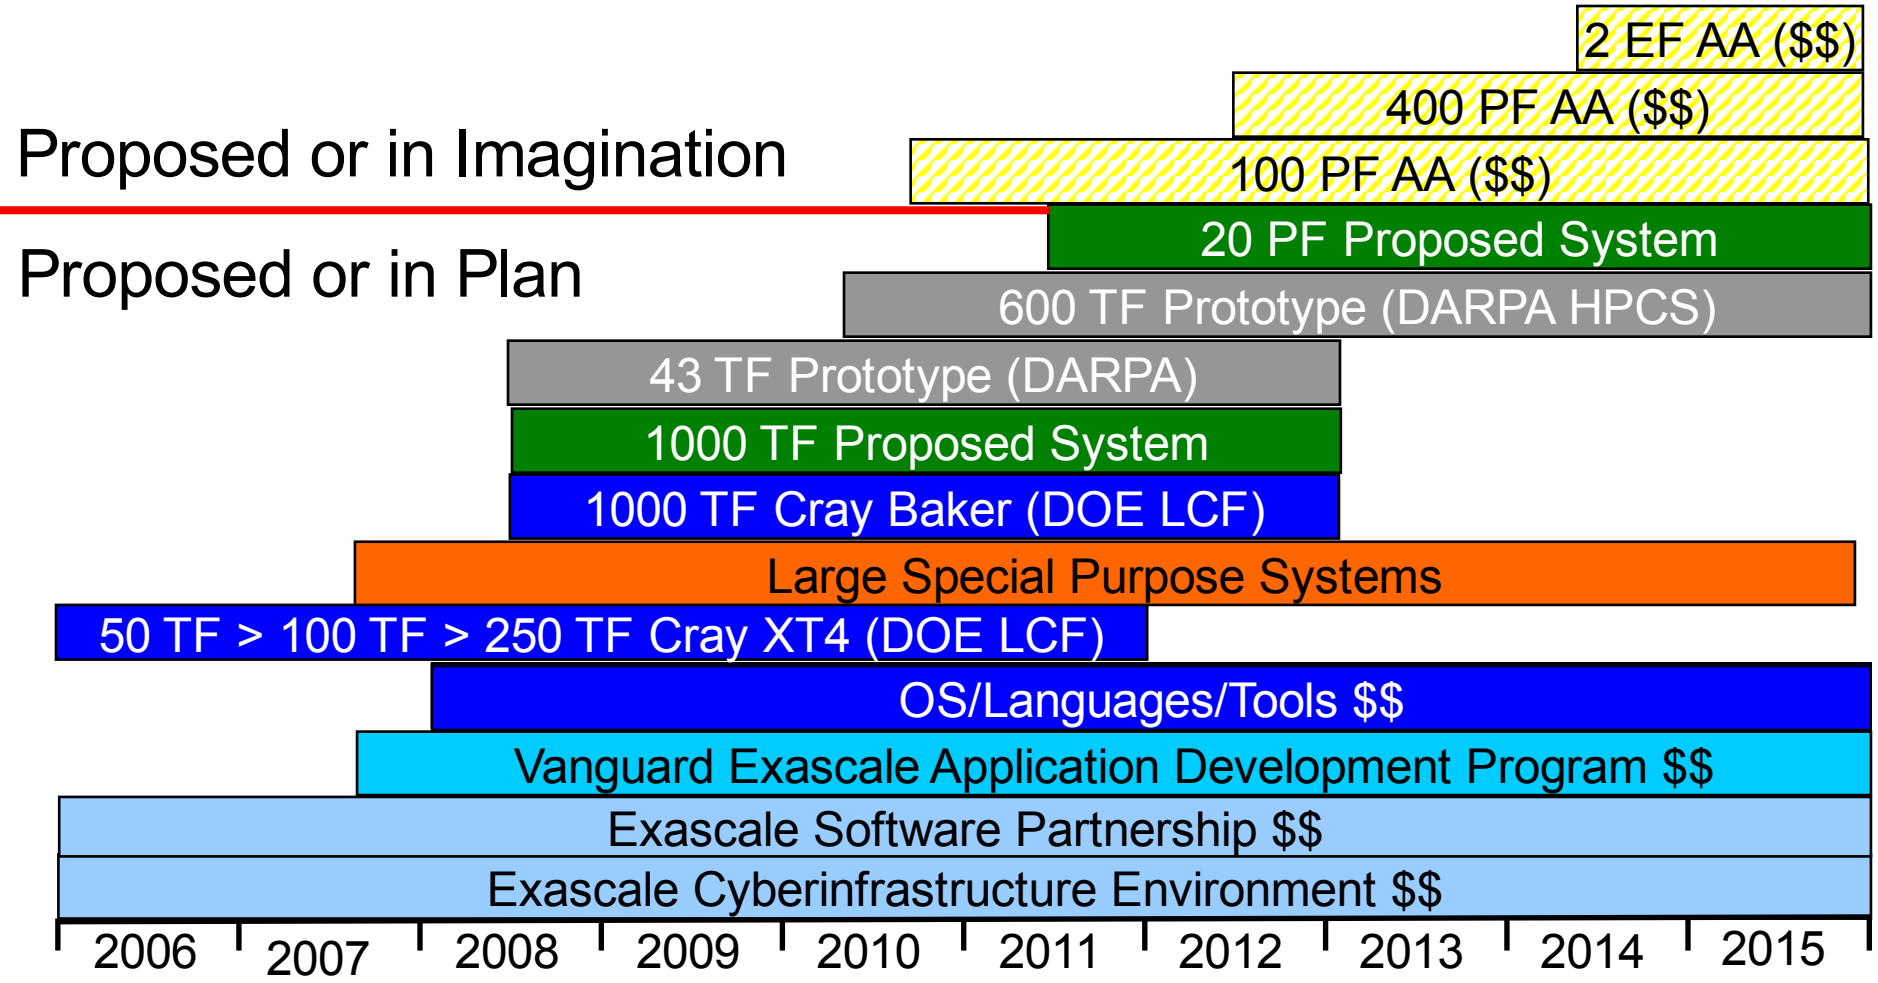

How to design a useful system:
Different problems, similar problem spaces, system impact the same

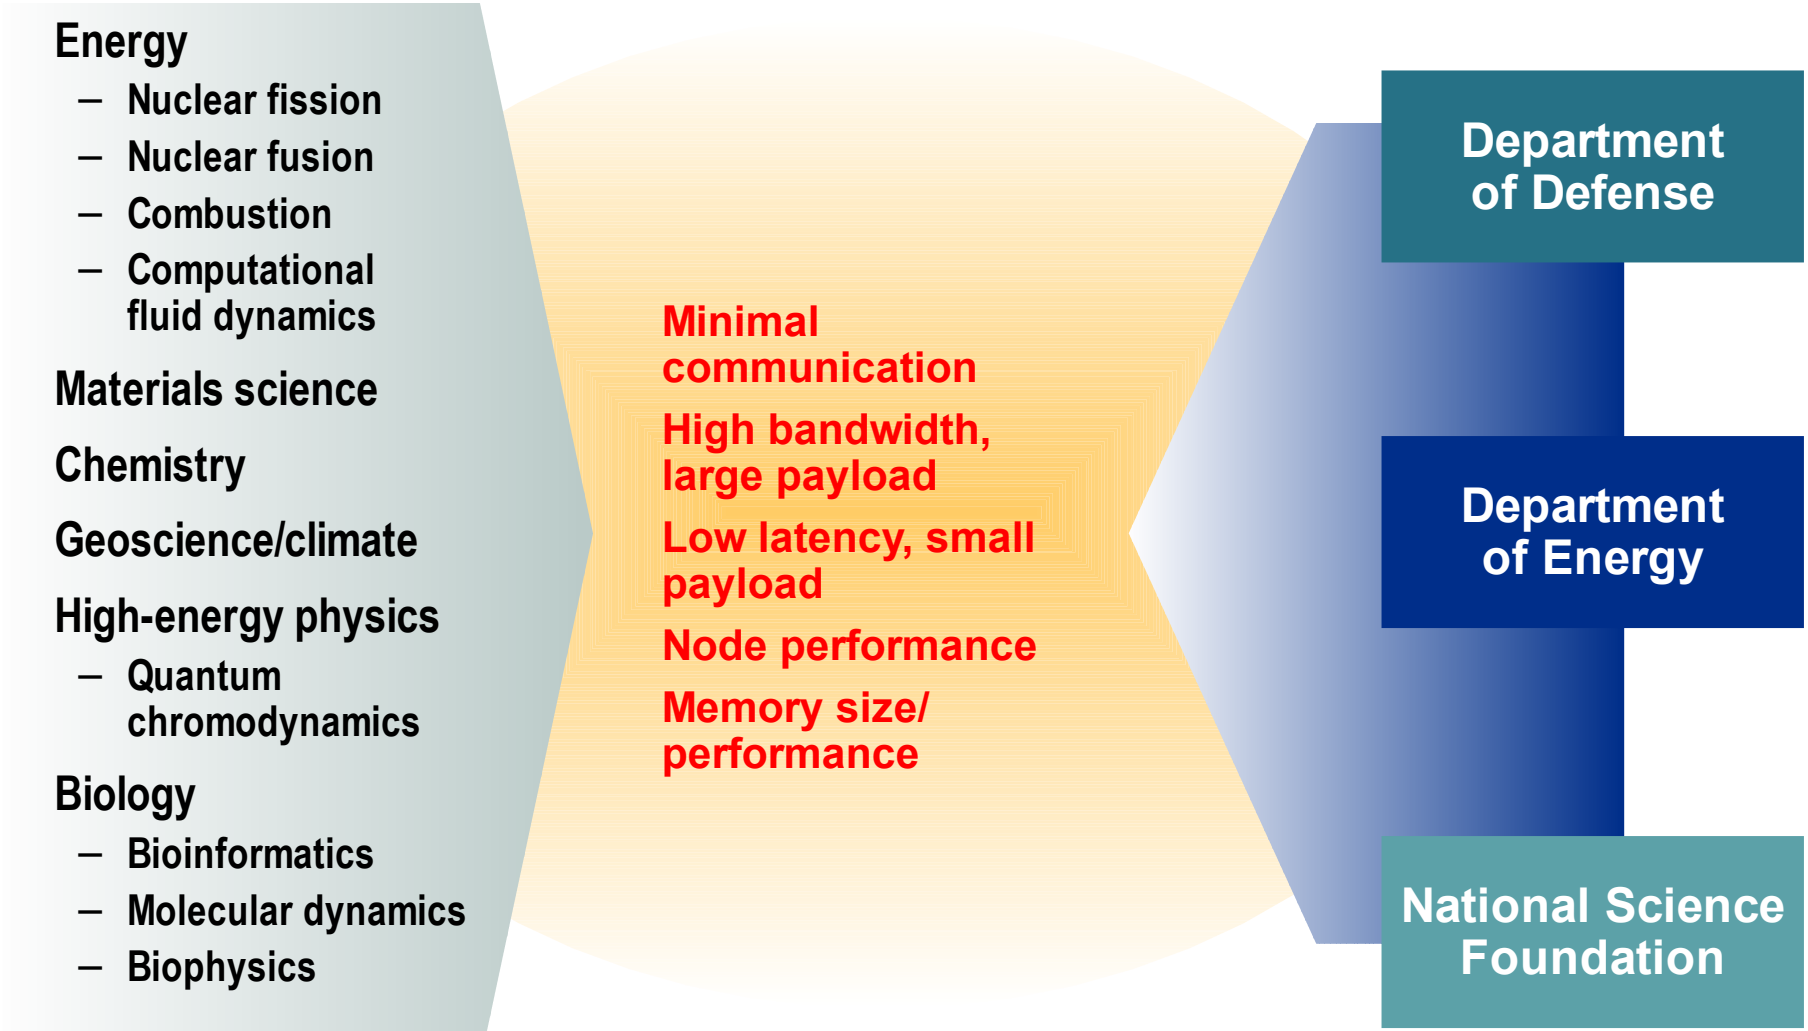

How to design a useful system:
Different problems, similar problem spaces, system impact the same

Energy
- Nuclear fission
- Nuclear fusion
- Combustion
- Computational fluid dynamics

Materials science

Chemistry

Geoscience/climate

High-energy physics
- Quantum chromodynamics

Biology
- Bioinformatics
- Molecular dynamics
- Biophysics

Highly multi-core

Co-Processors

Heterogeneous architectures

Additional level of memory hierarchy

Hybrid programming

Linear Algebra

Department of Defense

Department of Energy

National Science Foundation

System Types (Current & 2015)

Don't worry about DNA or QC (yet)

Infrastructure costs

BG/*,… (Lots-o-Processors, little memory)

FPGA, Cell, Cell follow on, GPU, GPGPU, PIM, Hybrid-(HW/SW)

Accelerators. Outside, then inside. History repeats itself

Power Series (P*) Complex, Power hungry,Maximum Integration

Cascade

XMT, Eldorado

Vector , Scalar , Threads , SPD … (Socket compatability)

We will most likely have more than one type in the future

We will need a variety of algorithms/implementations

We need to manage data, locality of reference

Most current systems limit the impact of our algorithms

Current systems limit the performance of our apps.

Sensor model

Without DT/Adv Engineering, we will see something like:

>80 M Cores (1.3M S, 64C/S)

130 MW

.000X B:F

Where do your algorithms fit in the scaling curve ?

Large Socket Solution

- DarkHorse / Pegasus (2Exa) (External forces $$ 2015)
  - ~10-40 TF/Socket (Hetero. Basis)
    - AMD + ATI (GPGPU) (“n”X86_64+”m”GPU/MVP)
    - Intel
    - IBM Cell+++ ?
    - FPGA , CAM layer
  - ~1-2 KW Socket (22nm <)
    - Coolable, but not necessary
      - Liquid Metal, Micro-channel cooling, spray cooling
      - Ceramic, Organic Substrate/Carrier
  - ~8-16G L3 - on die (3D)
  - ~50-400 PB Main (3D)
    - Self healing
    - S/G
    - Different Socket
    - >4B:F
  - ~400 GB/s node-to-node (All Optical, Proc+)
  - ~20PB/s global I/O BW
  - Full 128bit support (256 not yet)
  - MUST be balanced in Compute, Memory, Interconnect, Storage
  - Compiler / Data Flow ??
  - Some new technologies on the horizon

Cyclops / Intel (Polaris) (Multi-Cores are here)

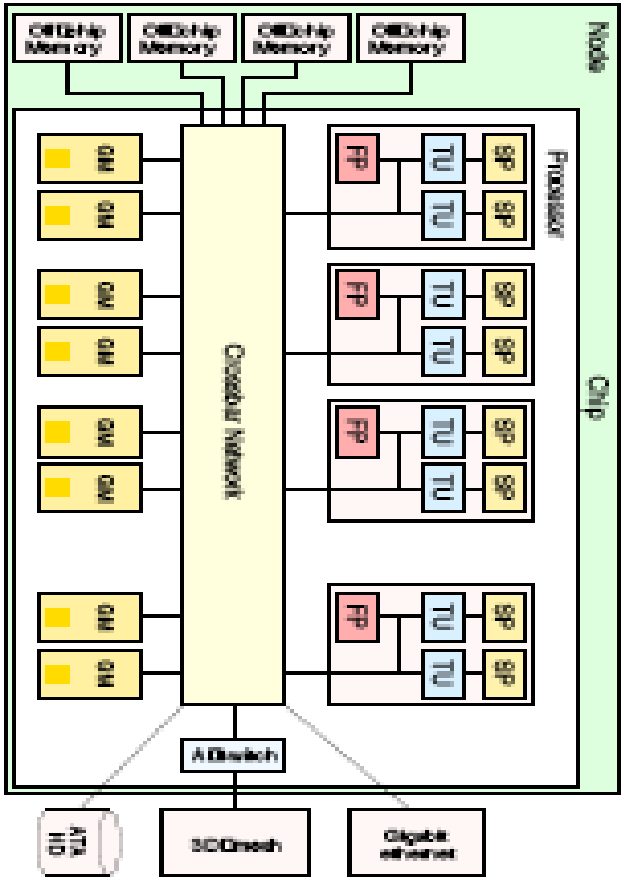

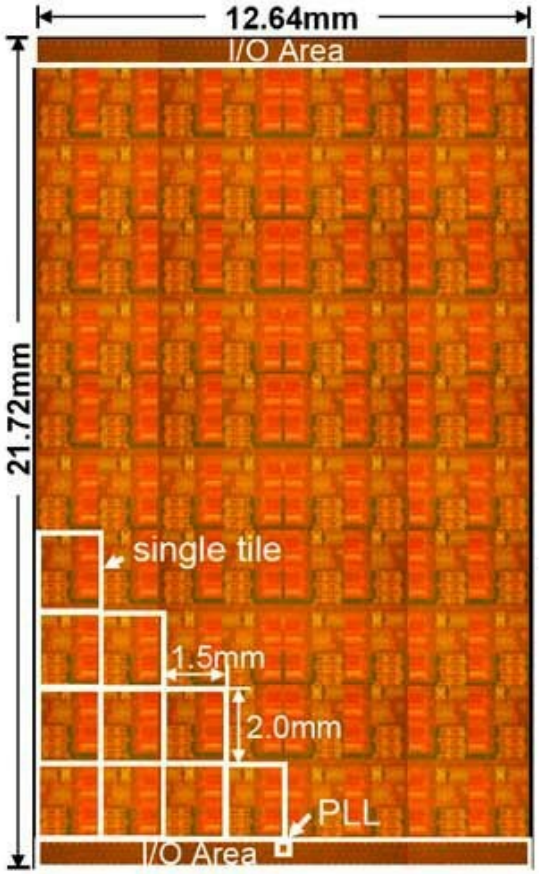

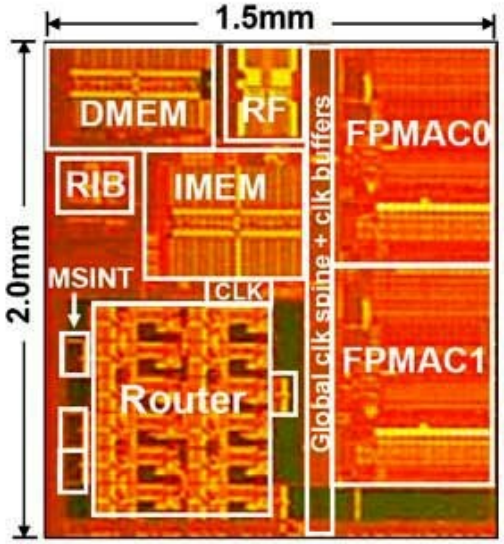

Technology	65nm CMOS Process
Interconnect	1 poly, 8 metal (Cu)
Transistors	100 Million
Die Area	275mm²
Tile area	3mm²
Package	1248 pin LGA, 14 layers, 343 signal pins

DarkHorse (A 3D System, 2002)

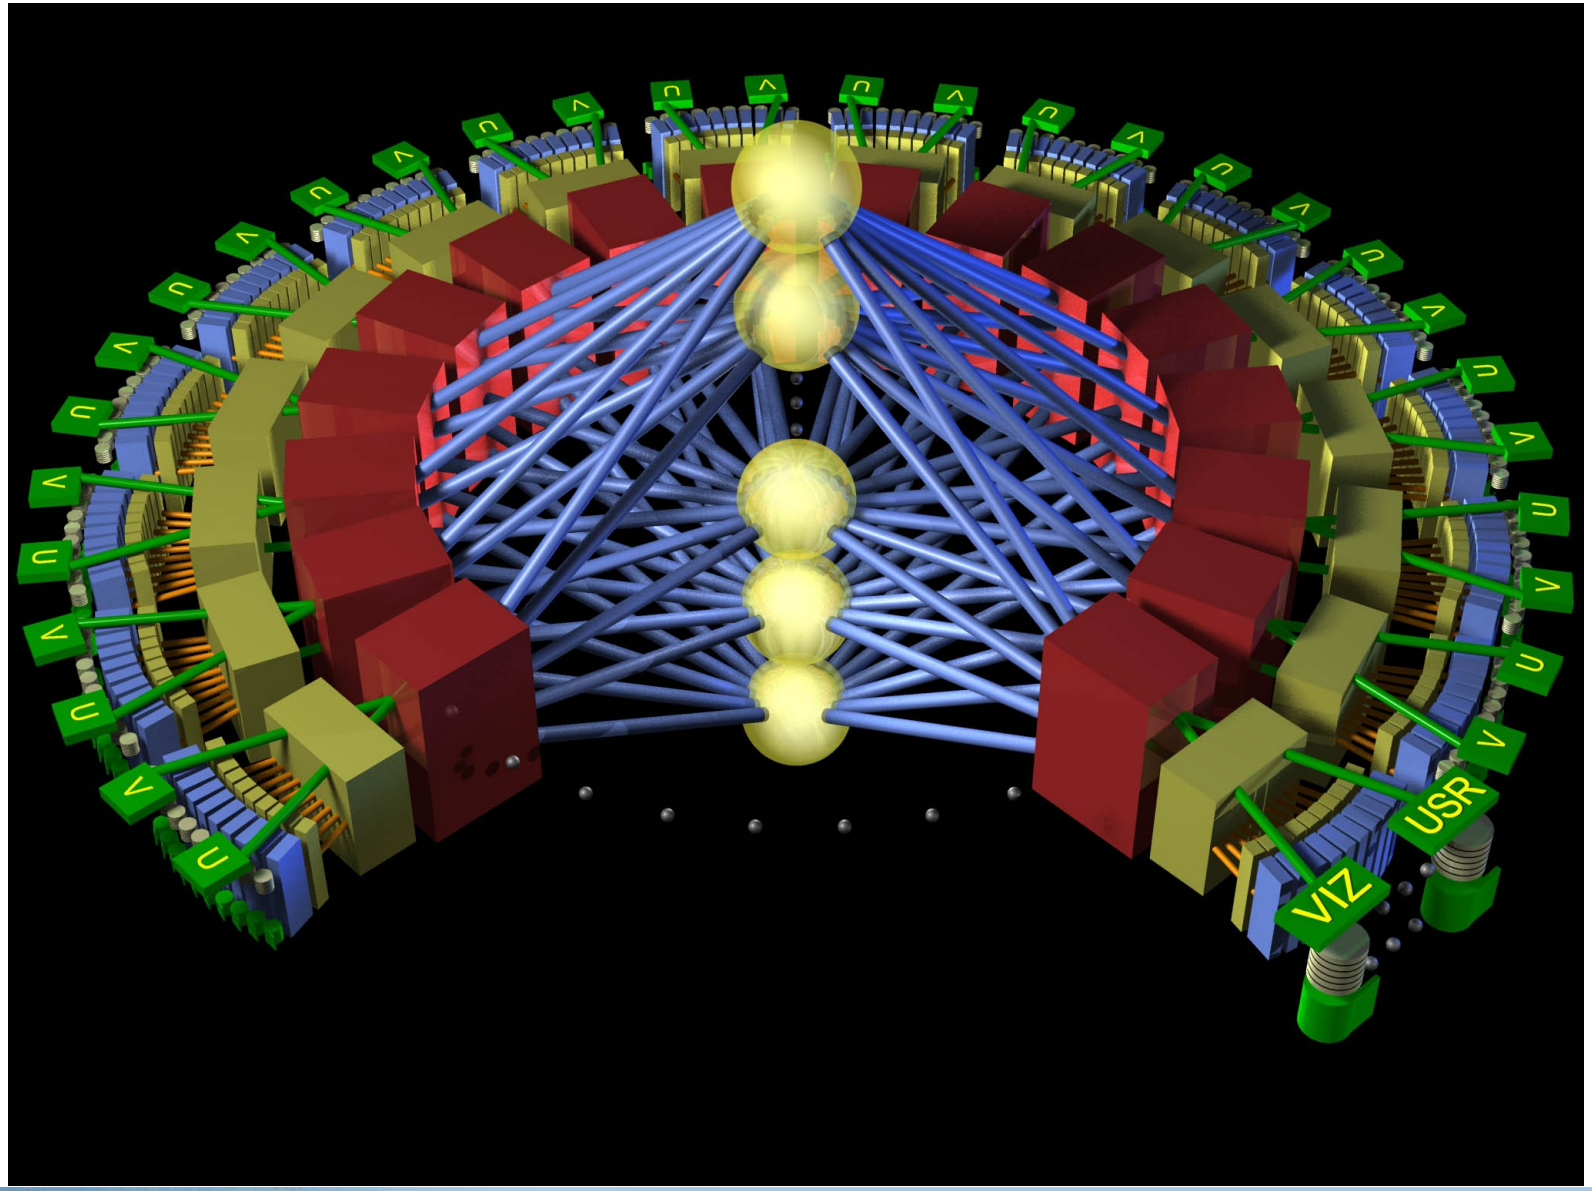

Conceptual Device Structure

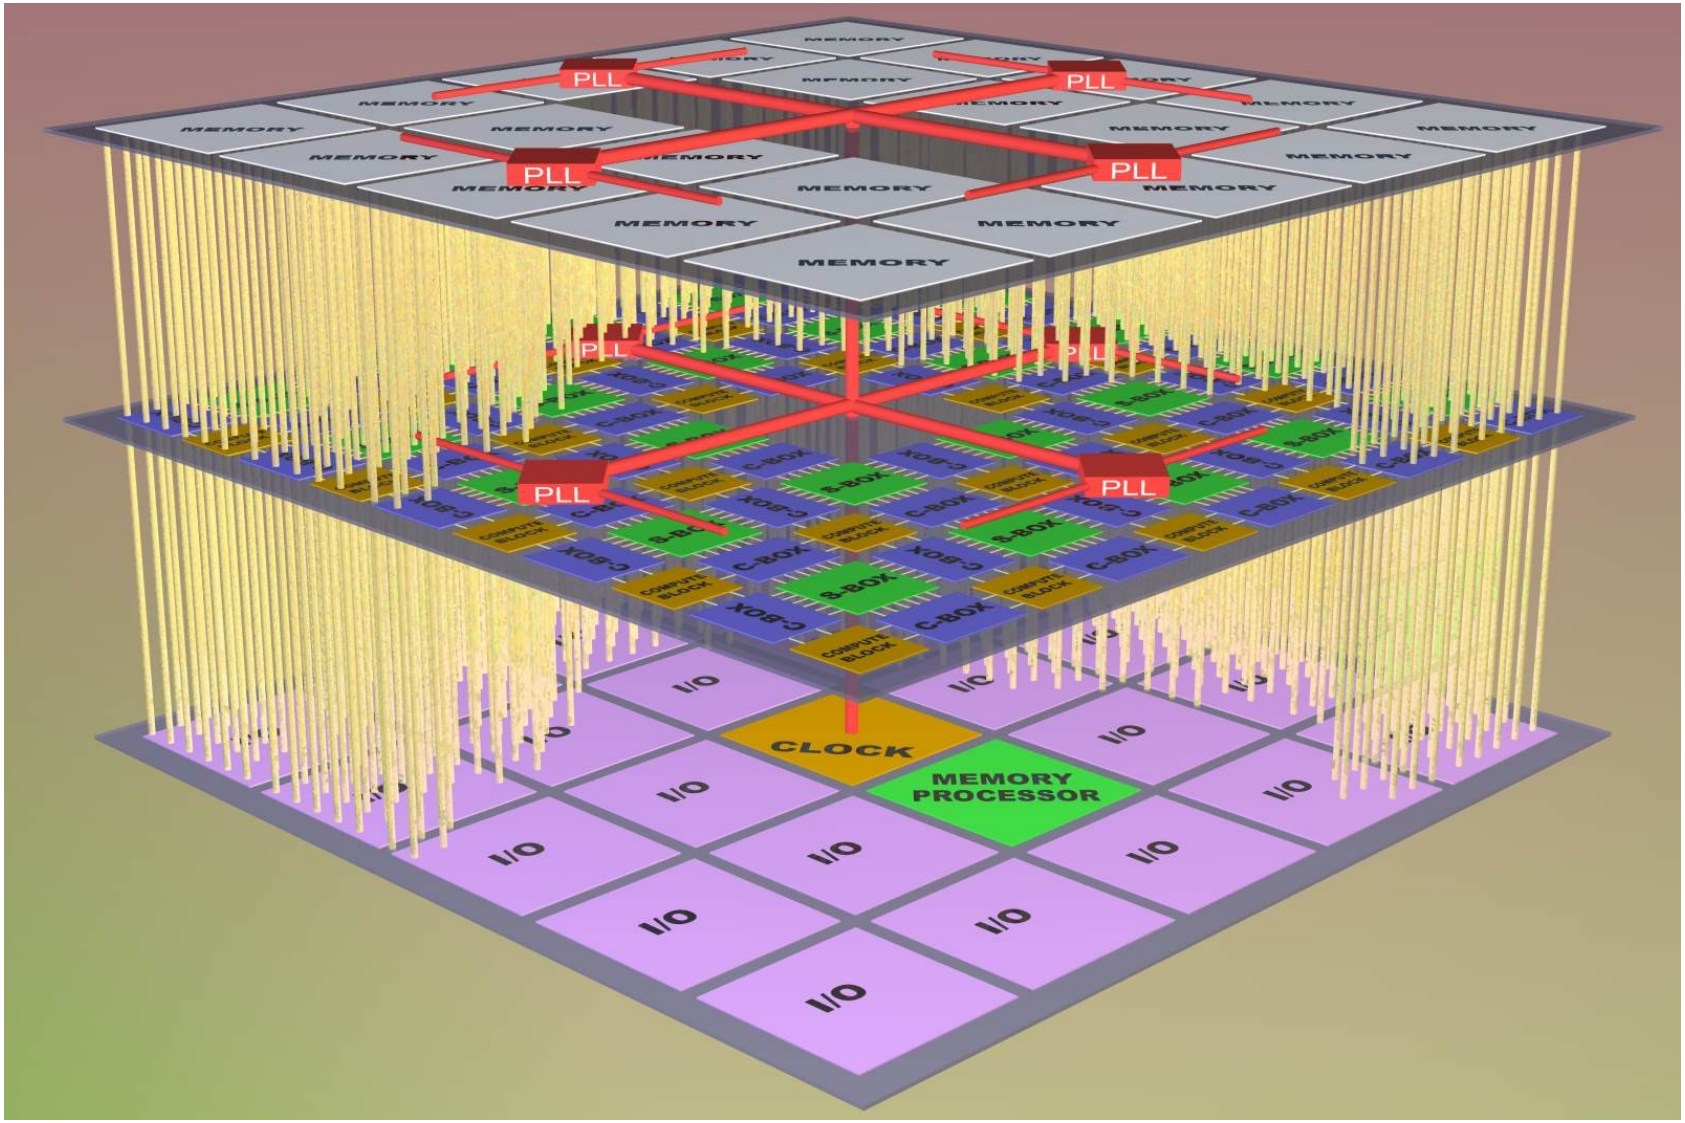

Progress (These were Disruptive Technologies)

3D (On Stack)
- ✓ Memory
- ✓ Memory + Processor , Memory + Sensor
- o Memory + Processor + FPGA + CAM + Network (Design only)

- 3D Network Chip (Off Stack)
- Optical Interconnect (C2C)
  - ✓ IBM Zurich (FR4 + Flex)
- Optical Switch
  - ✓ OSMOSIS
- Cooling
  - ✓ 1 - 2 KW
- Multi-Core / Multi-Thread
  - ✓ >64
- Hybrid / Heterogeneous
  - ✓ Cell (all other GPGPU's)
- We need to finish the development and commercialization process
- The 3D processes can be applied to many technologies.
- We need to continue to fund risky Research.

Backup