

These slides are at <http://www.studham.com/scott/presentations.php>

UT-BATTELLE



AMES LABORATORY



ARGONNE NATIONAL LABORATORY



Los Alamos NATIONAL LABORATORY



Lawrence Livermore National Laboratory



NASA



NCAR

Pacific Northwest National Laboratory



PPPL
PRINCETON PLASMA PHYSICS LABORATORY



Sandia National Laboratories



Panel Discussion
Future Technologies that may facilitate science breakthroughs

R. Scott Studham
Chief Technology Officer
National Leadership Computing Facility
Oak Ridge National Laboratory

March 21, 2005

THE CENTER FOR
COMPUTATIONAL SCIENCES

OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY

Future Technologies that may facilitate science breakthroughs

Multi Core CPUs
Multi Chip Vector Processors
Reconfigurable Computing Processors in Memory
New Chip/Memory interface
PS3: Streaming Vector
OS and filesystem advances
Co-Processors:
Physics Processing Units
Graphics Processing Units
Telco Parts
BLAS accelerators

For the sake of time I'm not going to cover.....

- Multi Core CPUs
- Multi Chip Vector Processors
- Reconfigurable Computing
- Processors in Memory

....and next, next generation stuff like
Optical and Quantum computers

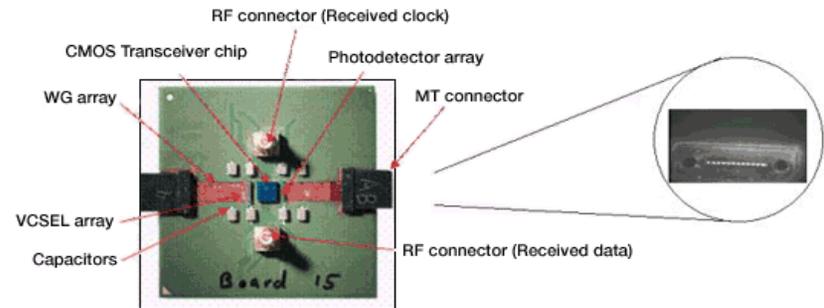
Change memory / processor interface

- Pins are slow / expensive.
 - Modest increase in frequency and number of Pins in not keeping pace with processor power.
 - Causing “Memory Wall”

- Some example solutions are:

- optical off the chip

- Non NDA information at:
 - Intel's chip-to-chip optical I/O interconnect technology
<http://www.intel.com/update/contents/it04041.htm>
 - “Proximity Communication”
 - Non NDA information at:
 - <http://research.sun.com/async/Publications/KPDisclosed/sml2003-0241/sml2003-0241.pdf>
 - Ivan Sutherland, “Face to Face Chips,” US Patent 6,500,696, Filed October 2, 2001, Granted December 31, 2002.



A completed prototype package for chip-to-chip optical data communication. The MT (multiterminal) connector is magnified to show the light output from the 12 channels of the transmitter.

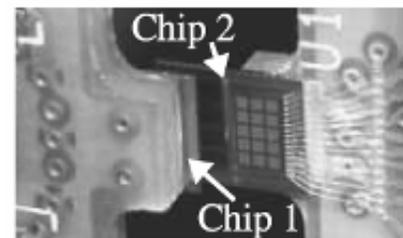


Figure 8. Detail of aligned chips

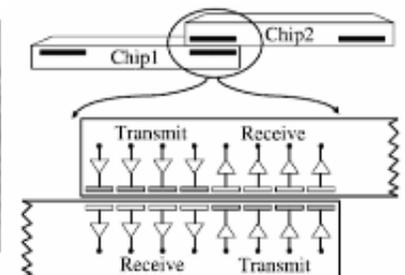


Figure 2. Cross-section of chips

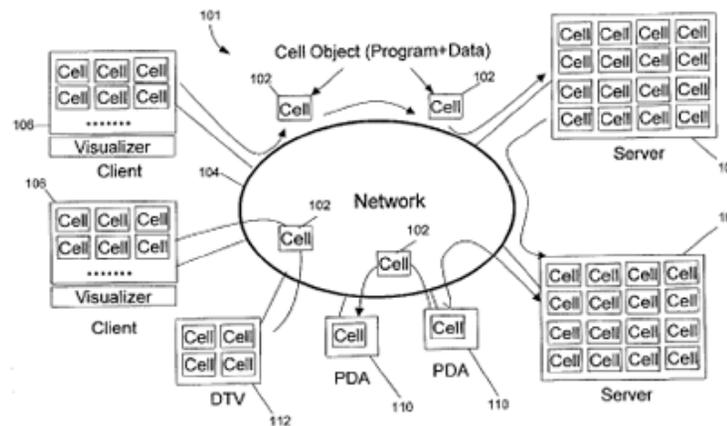
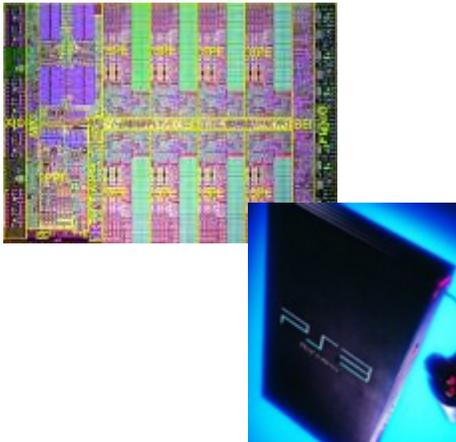
PS3: Streaming Vector

9 core “processor”

- one similar PowerPC G5 controller.
- Eight cores are called APUs and these are very high performance vector processors.
 - own block of high speed RAM
 - capable of 32 GigaFlops (32bit).
 - APUs are independent processors and can act alone or can be set up to process a stream of data with different APUs working on different stages.

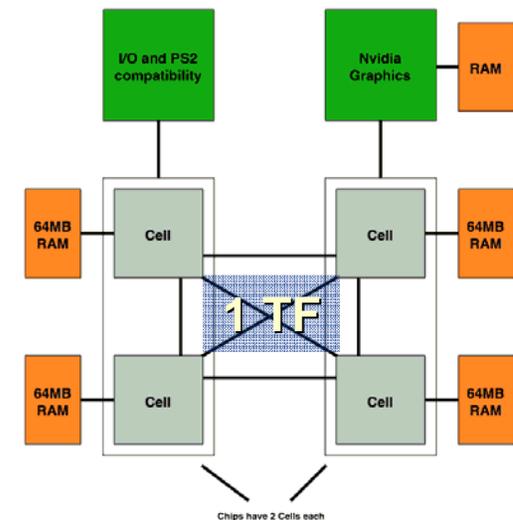
Cell:

- Clock speed over 4GHz.
- 100 GBytes per second aggregate Memory & I/O speed:
 - Dual XDR controller gives 25.6 GBytes per second.
 - Dual configurable interfaces give 76.8 GBytes per second.
 - Memory currently limited to 256 MB per Cell (PS3 likely limited to 64MB)
- Peak = 256 GigaFlops
- 221 square mm in 90nm.
- 234 million transistors
- IBM start manufacturing within 6 months, Sony to start later in the year.

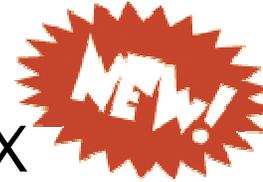


PlayStation 3 Architecture ?

This is a guess at what the PS3 architecture will look like.



Co-Processors



- Physics Processing Unit - AGEIA PhysX processor. PCIE, 128MB GDDR3, 25 watts, 1,000 faster than a PC at:

- Rigid body dynamics
- Universal collision detection
- Finite element analysis
- Soft body dynamics
- Fluid dynamics

“Walls and surroundings should be fully destructible within games - ever been driving a tank in a game like Call of Duty and been stopped by a shrub? Ageia's PPU is an important step in the right direction as it can take the current limit of 30-40 bodies of today's high-end CPUs to a maximum of 40,000.”

-- AGEIA Announcement of PhysX



- Graphics Processing Unit

- Attractive \$/FLOP.
- Showing up in CS papers at conferences (SC04) but hasn't made it to the mainstream applications yet.
- Still not sure of effective programming model.

Co-Processor (Cont)

- Telco parts
 - Broadcom BCM1x80: 4 processors per socket, MIPS64 instructions, 12.8GF, 14.4GB/s to memory, <23 Watts.
- BLAS accelerators
 - CSX600 is ~50 GFLOPS at <25 watts
 - 96 Processing Elements (PEs), 128 Kbytes of on-chip scratchpad SRAM
 - PCI-X (bad B/F ratio)



OS and Filesystem Advances

- The search capability at Google uses:
 - 200,000 processors
 - 2-4PB of disk storage
 - 15+ MW of power (my estimate)
 - A custom filesystem and OS.
 - Rob Pike (The Plan9 guy who B. Maccabe said “OS Research is dead”) is now at Google. OS Research isn’t dead, it just moved to industry.
- Amazon.com has a similar investment in hardware.
- Could you imagine the informatics problems we could solve if we were to create such a system?

Future Technologies that may facilitate science breakthroughs

	Usability	Market Cap	Challenge
Multi Core CPUs	High	Huge	Being designed for Oracle and Windows
Multi Chip Vector Processors	High	Small	Only a handful of companies left making these. If forced to choose, ISV's are focusing on scalar.
Reconfigurable Computing Processors in Memory	Low	Medium	None: If you are a EE PhD who does place and route
New Chip/Memory interface	??	??	Compiler support. Cost?
PS3: Streaming Vector	High	Medium	Patents and lawyers.
OS and filesystem advances	Low	Huge	Low memory address space, programming models.
Co-Processors:	High	Low	Rob Pike now makes a "Googlesalary". 100,000 CPUs
Physics Processing Units			Programming models.
Graphics Processing Units	Medium	Huge	Just coming to market.
Telco Parts	Low	Huge	Designed to send data to, not get data from.
BLAS accelerators	Medium	Medium	Slow parts. Small market for compiler support
	Medium	Small	No long term market.