

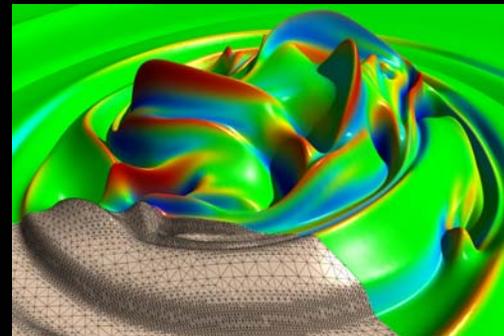
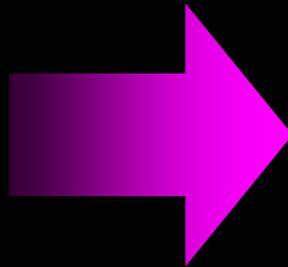


*Engineering Opportunities in the 21st Century*

*March 12, 2003*

# **Streaming Through My PlayStation:**

*Future Directions for Hardware and Software Architectures*



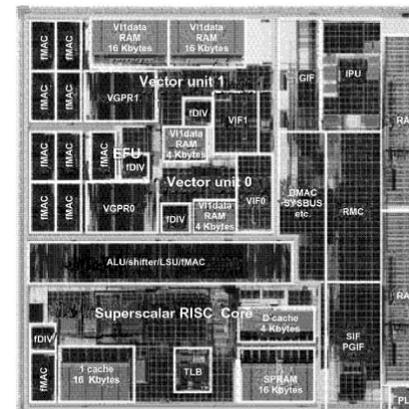
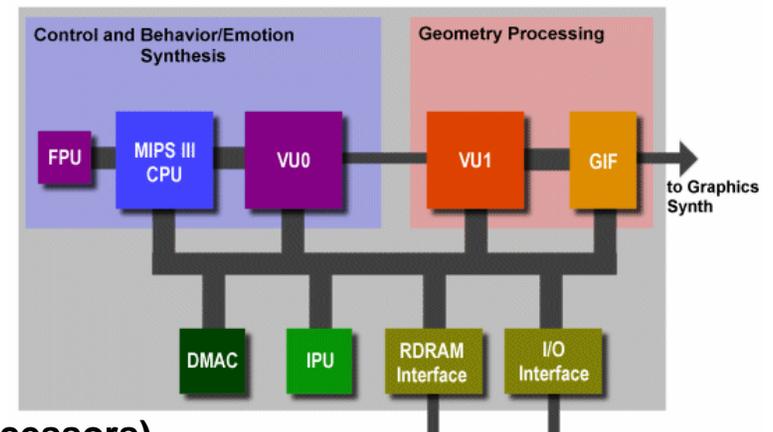
**Randall Frank, VIEWS Visualization Project Lead**

# The Sony Playstation



- Are graphics trends a glimpse of the future?
- The PlayStation2 architecture
  - 32MB RAM, 300Mhz, 6.2GFLOP
  - Streaming, custom serial streams
  - Cheaper to compute than to store
- The PlayStation3 (from the patent)
  - Core component is a cell
    - 1 PowerPC CPU + 8 APUs (“vectorial” processors)
    - 4GHz, 128K RAM, 256GFLOP/cell
    - Building block for multimedia framework
  - Multiple cells
    - Probably a four cell architecture (1TFLOP)
    - Central 64MB memory
    - Switched 1024 bit bus, optical links?

The Emotion Engine

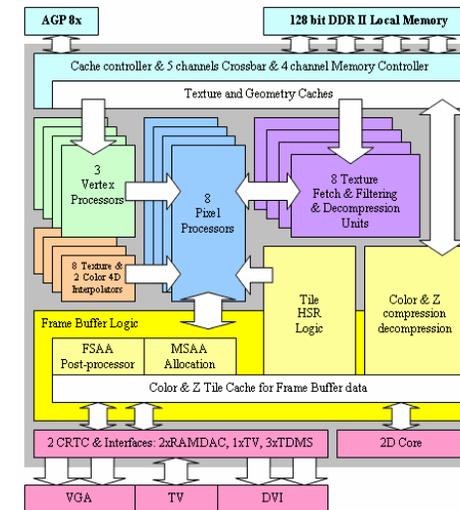


# Hardware Trends From Graphics



Engineering Opportunities in the 21st Century

- **Cheap FLOPS and MIPS**
  - GeForceFX: 96 IEEE 32bit units @ 500Mhz
  - Custom computation (e.g. MPEG)
- **CPU clock speeds soaring**
  - “Vector” clocks only slightly slower
- **Memory**
  - Not keeping pace with clock speeds
  - Data access specific caches (e.g. pixel tiles)
- **External pipeline bottlenecks (e.g. AGP, PCI, network)**
  - Multiple, custom paths between units
- **Micro-programmable hardware**
  - Fragment, vertex programs
  - Image and sound filters
  - Network stacks in NICs



# So what's new?



Engineering Opportunities in the 21st Century

- **Applications and standards**
  - **Multimedia: video and audio**
    - MPEG, DCTs, FFTs
  - **Gaming graphics**
    - Polygons, multi-texturing, vectors
  - **Common APIs**
    - Very focused functionality
      - OpenGL, Direct3D, DirectShow
    - Domain specific programming being exposed
- **There is a huge market for these systems**
  - Hardware is cheap and ubiquitous
  - Vendors looking to differentiate
- **Can we leverage this amazing horsepower, in a scalable fashion?**
  - Computational “streaming”



# The Streaming Programming model



- **Streaming exposes concurrency and latency at the system level as part of the programming target**
- **Data moves through the system: exposed concurrency**
  - Avoid global communication: prefer implicit models (e.g. Cr)
- **Memory model: exposed latency/bandwidth**
  - Scalable, must support very small footprints
  - Distributed, implicit flow between each operation
- **A working model:**
  - Computational elements + caching and bandwidth constraints
  - External “oracle” for system characterization and realization
- **Goals:**
  - Optimally trade off computation for critical bandwidth
  - Leverage traditionally “hidden” programmable elements

# Streaming Impacts on Software Design

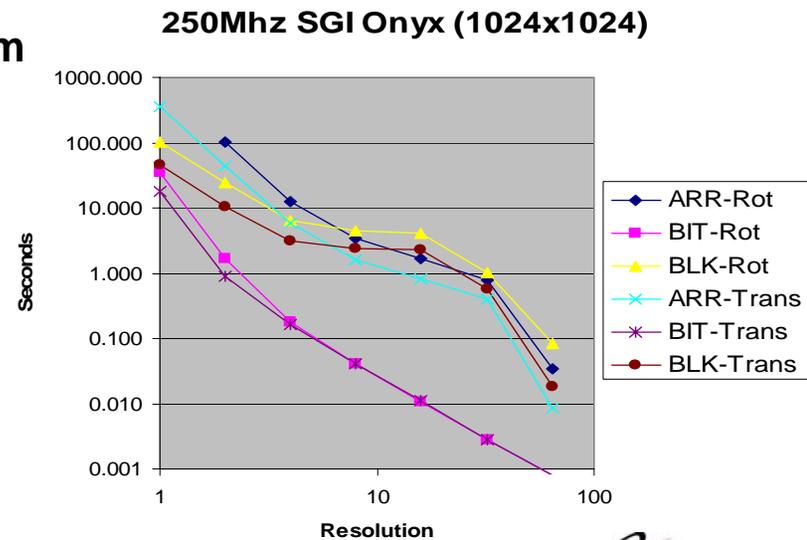
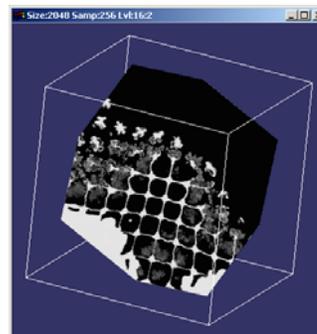
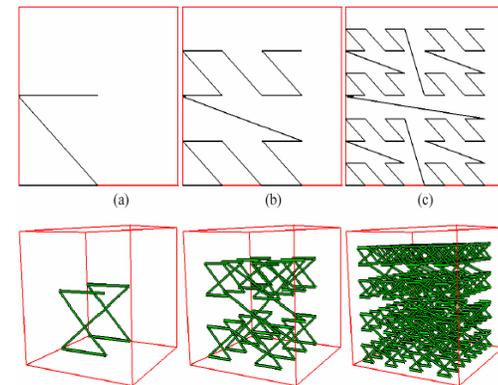


- **How does one target this model?**
  - **Integrated data structure and algorithm design**
    - Run-time targets
    - Algorithm remapping
  - **“Intent” expressive and architecture aware**
    - Abstract run-time compiled languages
  - **Memory design is key**
    - “Small” memory models, out-of-core design
    - “Cache oblivious” data flow
- **Implementations**
  - **New languages: Cg, Brook, DSP-C, Stream-C**
  - **Hidden beneath layers of API: OpenGL, Chromium, Lustre**
- **It sounds like a lot of work, what can really be gained?**

# Multiresolution Array Access From VISUS



- Arbitrary, multi-resolution array data access
- Integrated algorithm/data structure design
  - Data “reordering” on a spacefilling curve
    - In place transformation
    - Reordering is a simple bit manipulation
  - Cache oblivious
    - Arbitrary blocking is supported
  - Coupled asynchronous query system
    - Parallel rendering and queries
    - RAM used as a cache
- Example
  - Slicing 8B cells
  - 15MB RAM



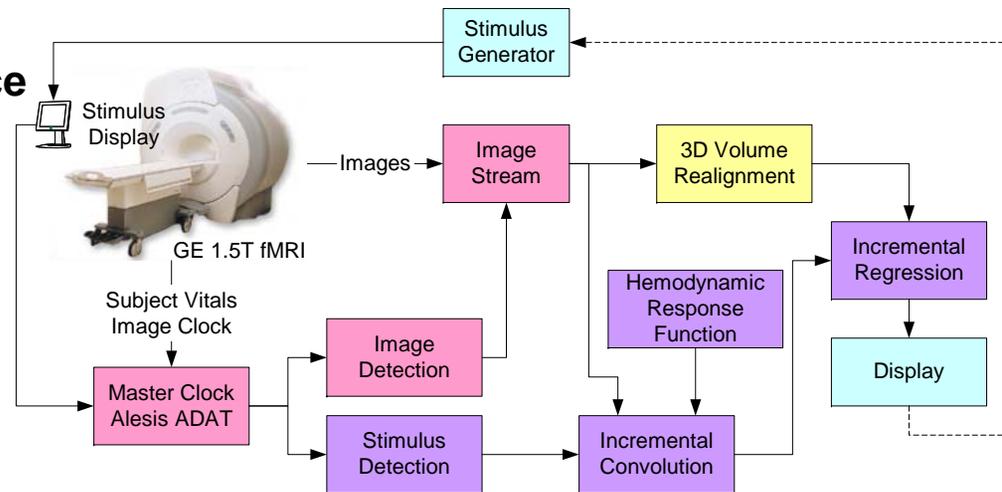
# IOW/A: Neuroanatomic Data Streaming



- **Work at the University of Iowa, Department of Neurology**
  - With: Brent Eaton, Lizann Bolinger, Thomas Grabowski
- **Goal: A system capable of real-time analysis of cognitive function.**
  - Allow direct investigation into individual's cognitive process
  - Interactive experiments and statistical testing of specific hypotheses

- **Design:**

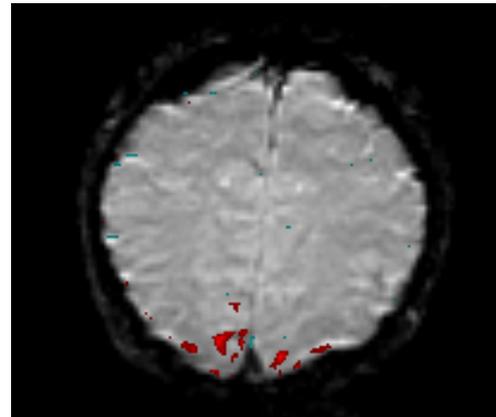
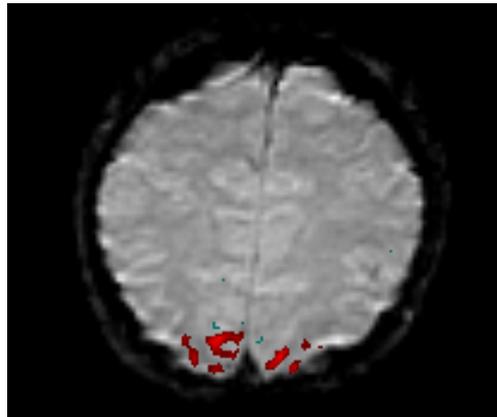
- Single data stream interface
- Arbitrary sample timing
- Fixed memory footprints
- Incremental algorithms
- Progressive alignment
- Multiple systems



# Insight Into The Working Human Mind



- **First complete real-time system of its type**
  - Performance envelope easily met by streaming (71ms/slice)
  - Lower resource requirements than previous off-line systems
    - A PC, a Sun and an SGI Octane vs a Cray T3E
  - Nearly optimal use of parallel, distributed resources
- **Unique features:**
  - Integrated 3D alignment (fixed temporal window)
  - Feedback for optimal acquisition or “Experimental Steering”



## Where does this lead?



- **Computation and memory caches everywhere**
  - PDAs, NICs, Drive controllers, Switches, Cell Phones
  - Potentially enormous “wall clock” performance gains
- **Utilizing them may require a disruptive shift**
  - New languages with higher levels of abstractions
  - Run-time “realization”, dynamic compilation and scheduling
  - Cache oblivious algorithms
- **Challenges with the opportunities**
  - Portability: how much work to retarget
  - Longevity: what are the implementation windows
  - Robustness: more points of failure
  - Debugging
    - More “conceptual” errors, but distributed systems & embedded systems

# Auspices



- **This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.**
- **This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.**
- **UCRL-PRES-152188**