



**SIGGRAPH2006**

SIGGRAPH2006

# A Performance-Oriented Data Parallel Virtual Machine for GPUs

Mark Peercy

Mark Segal

Derek Gerstmann



SIGGRAPH2006

# Problem Statement

“...significant barriers still exist for the developer who wishes to use the inexpensive power of commodity graphics hardware, whether for in-game simulation of physics or for conventional computational science. These chips are designed for and driven by video game development; the **programming model is unusual**, the **programming environment is tightly constrained**, and the underlying architectures are largely secret. The GPU developer must be an expert in computer graphics and its computational idioms to make effective use of the hardware, and still pitfalls abound...”

— **Course Description, SIGGRAPH  
2005 GPGPU Course**



# GPU as Compute Device

- **Interest for using GPU for compute**
  - Linear Algebra
  - Physics
  - Convolution
  - Sorting
  - FFT
  - Rendering (final-frame)
- **Exercises a small fraction of features in graphics hardware**



# Current GPU Abstraction

- **Rendering Pipeline** (*OpenGL, Direct3D*)
  - Graphics-centric programming model
    - Forced to manage graphics state
      - Even for *non-visual* computation(!)
  - Implemented through graphics driver
    - Mechanism designed to **hide** hardware
    - *Imposes critical policy decisions*
      - How + when + where data resides
      - Optimized for graphics and games
    - Driver updates exhibit different behavior

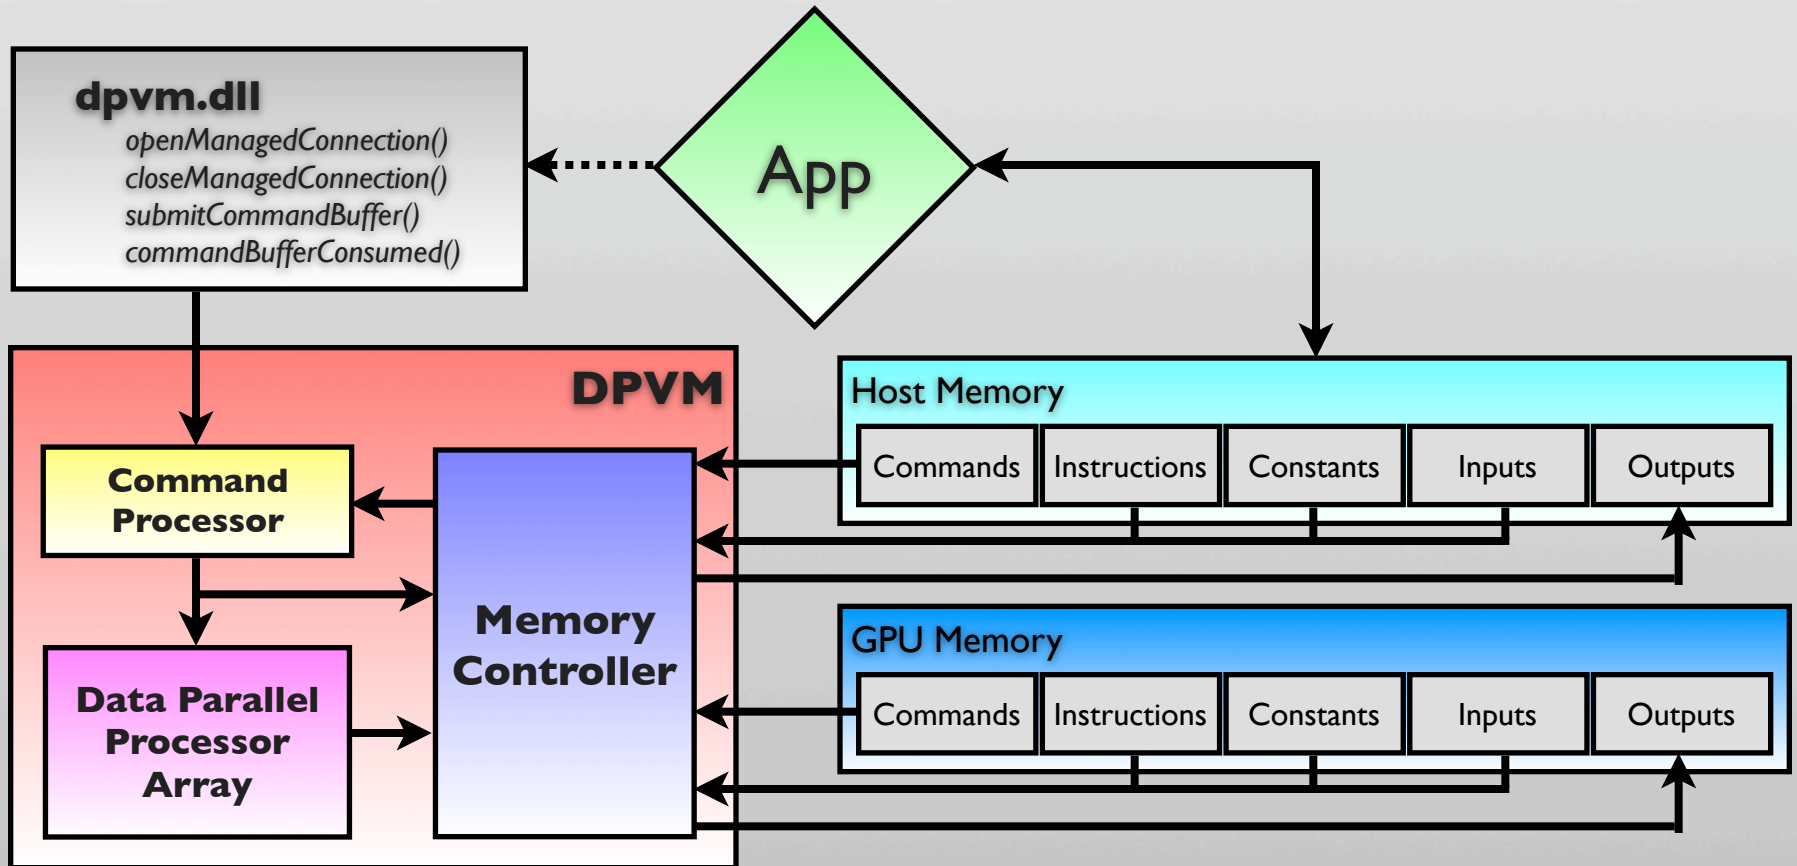


# A Data Parallel Approach

- **The Data Parallel Virtual Machine (DPVM)**
  - Expose relevant parts of the GPU as they **really** are
    - *Command Processor*
    - *Data parallel processor arrays*
    - *Memory controller*
  - Hide all other graphics-specific features
  - Provide direct communication to device
  - Eliminate driver implemented procedural API
    - Push policy decisions back to application



# The Data Parallel VM



# Command Processor

- **Abstracts communication from architecture**
  - Commands are architecturally *independent*
- **Accepts command buffers (CBs) in memory**
- **Interprets commands in buffer**
- **Distributes work to processor array**
- **Application manages command buffers**
  - Application fills and submits CB
  - Application handles synchronization





# Command Processor

- **Complete List of Data Parallel Commands**

## Program Execution

- `set_cond_val`
- `set_domain`
- `start_program`
- `set_out_mask`
- `set_cond_out_mask`
- `set_cond_test`
- `set_cond_loc`

## Cache Control

- `inv_inst_cache`
- `inv_constf_cache`
- `inv_consti_cache`
- `inv_constb_cache`
- `inv_cond_out_cache`
- `inv_inp_cache`
- `flush_out_cache`
- `flush_cond_out_cache`

## Memory Layout

- `set_inst_fmt`
- `set_inp_fmt`
- `set_out_fmt`
- `set_cond_out_fmt`
- `set_constf_fmt`
- `set_consti_fmt`
- `set_constb_fmt`

## Performance Counters

- `init_perf_counters`
- `start_perf_counters`
- `stop_perf_counters`
- `read_perf_counters`



# Data Parallel Processors

- **Performs floating-point computations**
- **Accepts binary executable (ELF)**
  - Formal application binary interface (ABI)
  - Uses **native** instruction set architecture of processors (ISA)
    - ISA is architecturally **dependent**
    - Only ISA needs to be updated for new architectures
- **Application submits compiled binary**
  - Executable is immune to driver changes



# Memory Controller

- **Services GPU requests to read/write memory**
  - Exports graphics memory directly
    - GPU memory (accessible by GPU only)
    - Host memory (accessible by CPU + GPU)
  - Application manages memory
    - Specifies locations and formats
    - Can **cast** between formats (w/o copying)
    - Controls cache invalidation



# ATI X1k DPVM

- **Implementation** (*CTM - Close to the Metal*)
  - Radeon X1k architecture (eg X1900)
    - Exposes hardware resources (SM3.0 DX9+)
    - Native ISA (ASM Text + Binary Formats)
  - Runtime library
    - Low-level driver components
  - Support libraries
    - Command buffer packing
    - Memory allocation
    - Assembler/Disassembler



# ATI X1k DPVM

- **Processor Resources** (*ATI Radeon X1900*)

<b>x16 Inputs (textures)</b>
• float1/2/4
<b>x4 Outputs (MRT) ...</b>
• float1/2/4
• assigned (x,y)
<b>... or xINF Outputs</b>
• float1
• arbitrary (x,y)
<b>x512 Instructions</b>
• any combination...
ALU / FLOW CONTROL
INPUT / OUTPUT










<b>x256 Float Constants</b>
• float4
<b>x32 Integer Constants</b>
• int4
<b>x32 Boolean Constants</b>
• bool1
<b>x128 Registers (GPR)</b>
• float4



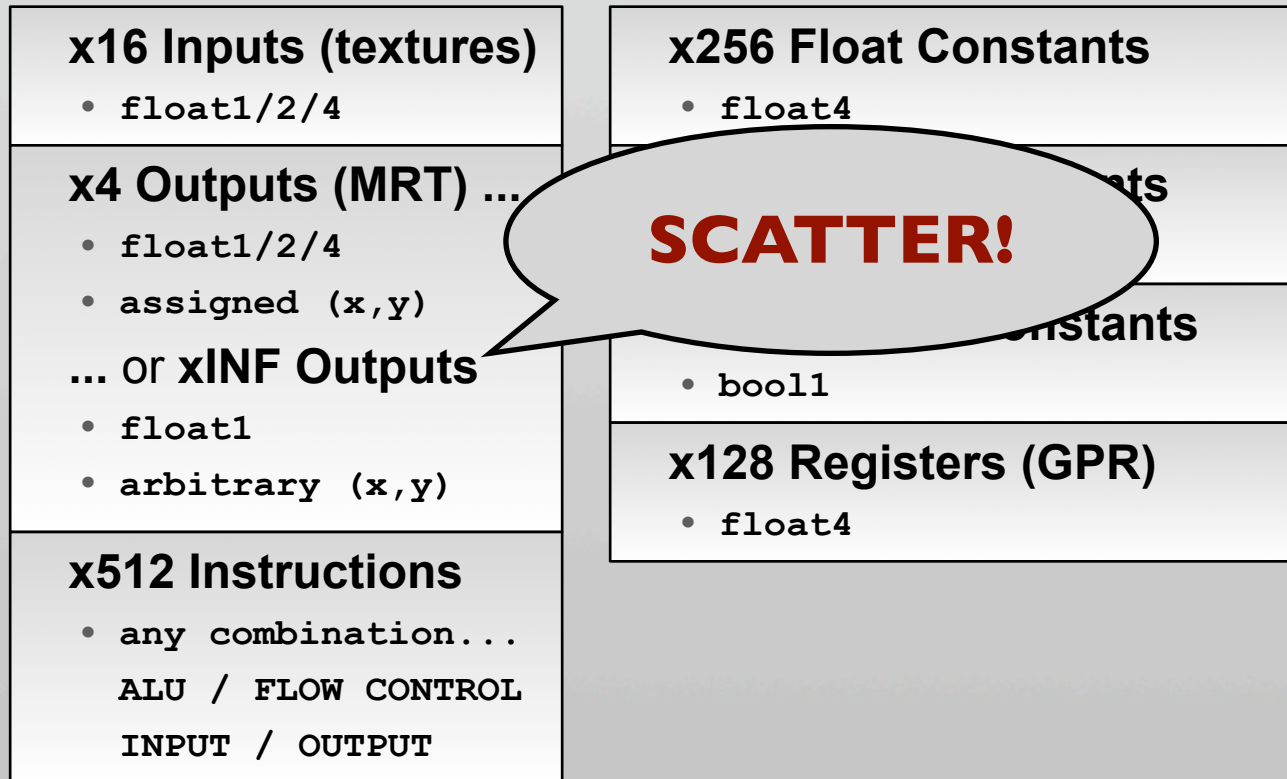







# ATI X1k DPVM

- **Processor Resources** (*ATI Radeon X1900*)



# ATI X1k DPVM

- **Additional Features** (*beyond SM3.0*)
  - **Scatter** (output `float` values to arbitrary locations)
  - Tiled memory formats
    - **Fetch4** (retrieve `x4 float` values in a single clock)
  - ABI w/native **ISA** allows hand-optimizations
  - Ability to read/write **directly** to/from **host** memory
  - Avoid non-IEEE floating point optimizations
  - Application dictates granularity of CB submission
  - Unlimited application execution time (arbitrary CB)



# ATI X1k DPVM

...DEMOS...



SIGGRAPH2006



# DPVM Applications

App	Benefit	Features
Matrix-Matrix Multiply	x10	CB, ISA, mem-formats, mem-offsets, interleaving, fetch4
FFT	x2	CB, ISA, interleaving
GPURay	x2	CB, mem-formats
QJulia	x2	CB, mem-formats

# Conclusion

- **Benefits of the Data Parallel Approach**
  - Straight-forward programming model
    - Allows hand-tuned optimizations
  - Exposes actual hardware device
    - Direct control over memory + processors
    - Application binary interface + native ISA
  - Application is responsible for all **policy** decisions
  - Allows **consistent** performance for compute



# Future Work

- **Other things to explore...**
  - Open area for tool development
    - Debuggers(!)
    - Statistical runtime profilers
  - New opportunities for compiler research
    - Support for high-level languages
    - Non-graphics optimizations



# Special Thanks...

- **ATI Research, Inc.**
  - *Mark Peercy, Mark Segal, Alpana Kaulgud, Raja Koduri, and everyone else...*
- **Stanford University**
  - *Mike Houston, Daniel Horn...*



# Questions?

**For more info contact:**

**[researcher@ati.com](mailto:researcher@ati.com)**



SIGGRAPH2006