

CPU And GPU Threading

- Many performance issue are the results of inefficient threading between CPU cores and/or GPU engines
- It is often hard to realize these problems exist using conventional performance tools
- This is an area where Event Tracing for Windows (ETW) can provide critical information in understanding performance issues

WinHEC
2008

ETW Traces


- ETW traces (file with .ETL extension) are generated using the XPerf tools
- ETL file may be viewed using various tools
 - XPerf - Powerful general purpose analysis tools
 - GPUView - Timeline view of a trace with emphasis on CPU/GPU interaction. Will be available with Windows 7 Beta SDK

WinHEC
2008

GPUView Is One More Tool

- GPUView doesn't replace F1, IceCap, SPU, Vtune, PIX, GPU PerfStudio, NVPerfHUD...
- GPUView's strength is in CPU/CPU, CPU/GPU and GPU/GPU interaction
- GPUView makes it easy to see if the CPU core and GPU engines are being threaded efficiently and if so where to look next to improve performance

WinHEC
2008




Threading Basics

```
Main()
{
    for (i=0; i<1000, i++)
    {
        DoSomething()
    }
}
```




Instrumentation

```
Main()
{
    for (i=0; i<1000, i++)
    {
        DoSomething()
    }
}
```



Sampling Analysis

```
Main()
{
    for (i=0; i<1000, i++)
    {
        DoSomething()
    }
}
```



WinHEC
2008


CPU Optimization

- We want to reduce the time it takes to execute the task. We care about the absolute time between Start and End
- F1, IceCap, SPU, Vtune allow us to drill down into the problem
 - Change algorithm
 - Use better instructions
 - Reordering loops
 - Be more mindful of cache misses
 - ...

WinHEC
2008

Optimized Thread

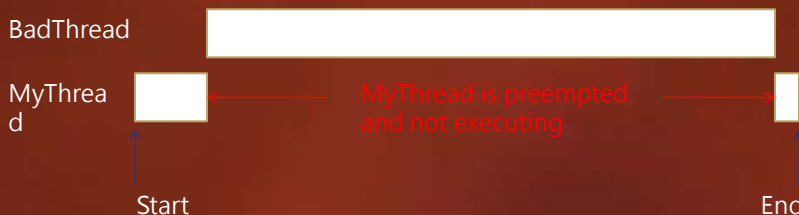
```
Main()
{
    for (i=0; i<100, i++)
    {
        DoSomething()
    }
}
```



WinHEC
2008

Multi-tasking Environment

- The threads being monitored for performance aren't the only threads running on the system other threads or driver (ISR/DPC) might interfere



WinHEC
2008

Device Interaction

- A thread might interact with other threads or with a device and may be put into a wait state waiting for something to happen on another thread or device




GPUView Process View



Idle Thread

- Every CPU core executes a particular thread at any given time; When Idle, a core executes the idle thread



WinHEC
2008

CPU Core

- Every CPU core gets assigned a color, threads are colored from the core they are executing on



WinHEC
2008

Thread Priority

- The priority at which a thread segment is executing is displayed in the left portion of that segment; If the segment isn't large enough, zoom in




WinHEC
2008

Thread Execution

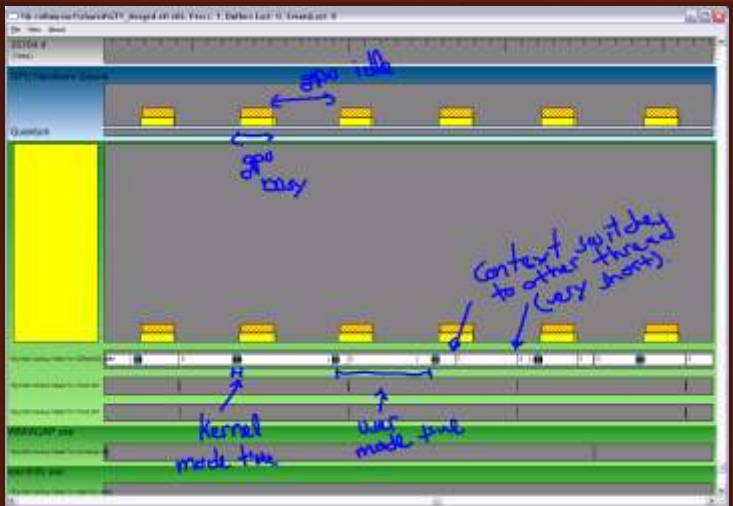
- Thread segments are colored to help identify what the thread is currently doing
 - Light blue: Thread in kernel mode
(only if syscall events were enabled as tracing option, not available on MP system)
 - Dark blue: Thread in dxgkrnl
 - Red: Thread in KMD



WinHEC
2008



Example: CPU Bound Game



GPU Context And Engine Queue



Context Queue

- The device queue (also referred to as the software queue or QueuePacket queue) represents the accumulation of work submitted by a process to a particular GPU Context
- A process can have multiple GPU contexts (ex: Multi D3D Device, LDA, Video...)

Context Queue

- Queue is represented through time as a histogram
- On submission of work by the UMD the histogram grows up, on completion of work by the GPU the histogram goes down



WinHEC
2008


Queue Packet

- Various type of packet may be queued
 - Orange: MMIO flip
 - Diag-Cross Hatch: Present through DMA
 - Cross Hatch: Deferred Command
 - Light Pink: WaitForSynchronizationObject
 - Dark Pink: SignalSynchronizationObject

WinHEC
2008

MMIO Flip


- Completion of the packet indicate the flip was pushed to the flip queue
- Flips are dequeued from the flip queue:
 - At VSYNC for synchronized flip
 - Immediately for immediate flip
- The content of the flip queue isn't currently displayed by GPUView
- Packet contain two pieces of information
 - VidPnSource : PrimaryAllocationHandle



WinHEC
2008

Present Through DMA

- Three types of present through DMA
 - DWM Redirection
 - Blit to back buffer
 - Immediate flip (through DMA)



WinHEC
2008

Deferred Command

- Mechanism used by VidMm to schedule work synchronously with a context execution; Ex:
 - Handle complex lock
 - Close allocation
- You don't usually have to worry about those

WinHEC
2008

Synchronization Primitive

- Packet inserted when UMD call Wait/SignalSynchronizationObject
- Packet contain that allow pairing of Wait/Signal; Signal of A3 will unwait Wait of A3

WinHEC
2008

GPU Engine Queue

- Also referred as hardware queue or DMAPacket queue
- Represent work (DMA Buffer) submitted to a particular GPU engine for execution
- Each exposed engine has exactly ONE engine queue



WinHEC
2008


DMA Packet

- Various type of DMA packet may be submitted to the GPU
 - Red: Paging
 - Black: Preemption
 - Diag-cross hatch: Present (through DMA)
 - Other: DMA payload from QueuePacket

WinHEC
2008

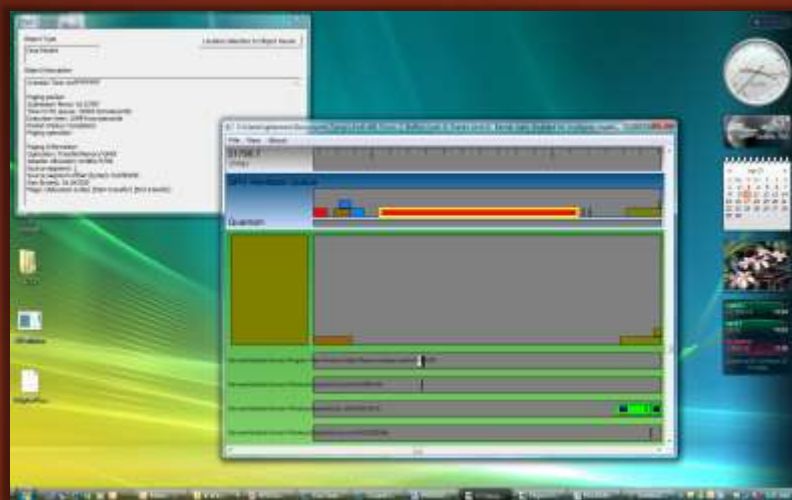
Paging Buffer

- Submitted as a result of a paging operation
- Cause of paging operation is usually the preparation of a DMA buffer; Look at the DMA packet following the paging operation
- Some exception
 - Pinning allocation
 - Evict allocation as a result of a DxLock
 - ...



WinHEC
2008


Paging Details



WinHEC
2008

Preemption

- Represent request by the schedule to preempt the DMA packet currently in the engine queue
- Preemption request don't have actual payload; Their completion indicate GPU acknowledgement of the preemption request




WinHEC
2008

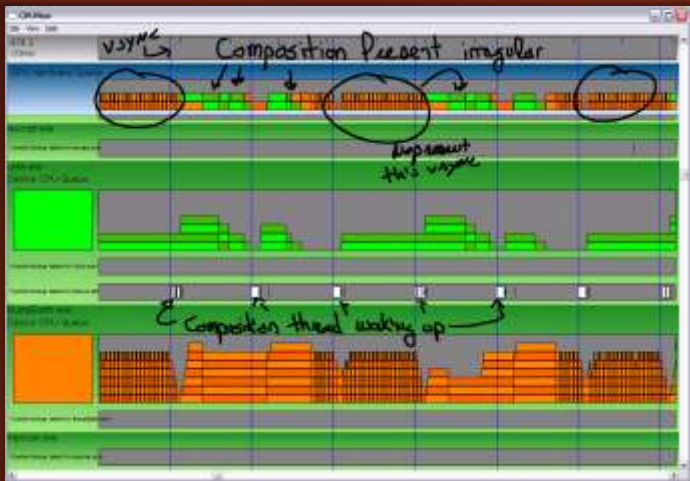
QueuePacket With DMA Payload



WinHEC
2008



Ex: DWM Regular GPU Priority



[illegible][illegible]

Call To Action

- The many core era is coming, add multi-threading support to your application to take advantage of new CPUs
- Use Xperf and GPUView along with your current set of performance tools to better understand your threading efficiency and enhanced the performance of your Windows application

WinHEC
2008

Additional Resources

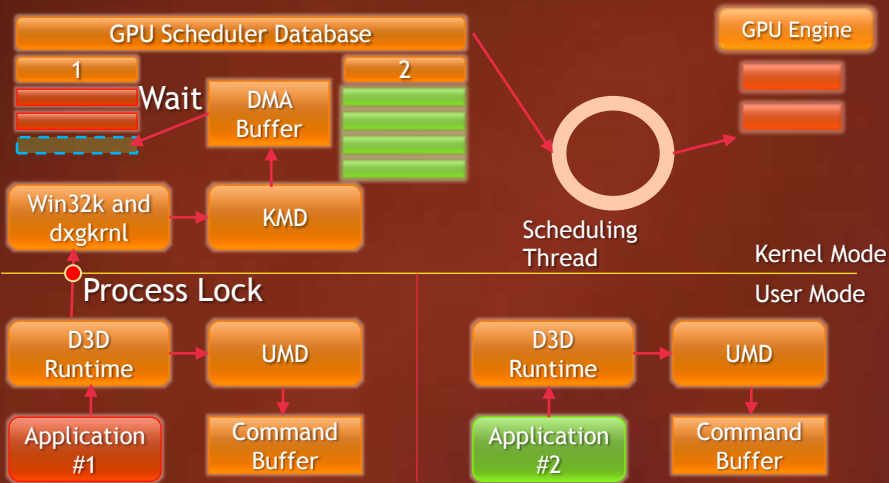
- XPerf
 - Available now as the WPT Kit which is part of the Windows 7 SDK
 - Older version also available online at
 - <http://www.microsoft.com/whdc/system/sysperf/perftools.msp>
 - Xperf documentation very useful to get started generating and analyzing ETW traces
 - <http://msdn.microsoft.com/en-us/library/aa139734.aspx>
- GPUView
 - Will be made available in Windows 7 Beta SDK as part of the WPT Kit
- For feedback and questions on GPUView use the following alias
 - gpuviewf@microsoft.com

WinHEC
2008

Appendix

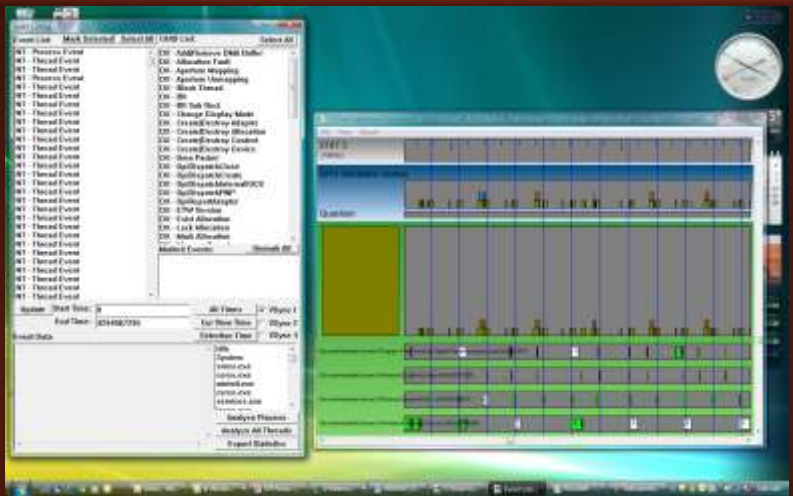
WinHEC
2008

WDDM GPU Scheduling




WinHEC
2008

VSync



WinHEC 2008

AllocationList And Placement




WinHEC 2008

AllocationList And Placement

- 0x83115AC0 (w) P0 Mem 0x87000 (540KB)
- Handle: Global identifier for the allocation
- Operation: (w) indicate a write operation. The default is read
- Relative priority: Where the allocation was placed relative to its priority
- Mem/App: Memory segment or Aperture Segment
- Size of allocation

WinHEC
2008

More Details On Allocation



WinHEC
2008

