GPU Virtualization on VMware's Hosted I/O Architecture

Micah Dowty, Jeremy Sugerman

VMware, Inc.

3401 Hillview Ave, Palo Alto, CA 94304

`micah@vmware.com, yoel@vmware.com`

Abstract

Modern graphics co-processors (GPUs) can produce high fidelity images several orders of magnitude faster than general purpose CPUs, and this performance expectation is rapidly becoming ubiquitous in personal computers. Despite this, GPU virtualization is a nascent field of research. This paper introduces a taxonomy of strategies for GPU virtualization and describes in detail the specific GPU virtualization architecture developed for VMware's hosted products (VMware Workstation and VMware Fusion).

We analyze the performance of our GPU virtualization with a combination of applications and microbenchmarks. We also compare against software rendering, the GPU virtualization in Parallels Desktop 3.0, and the native GPU. We find that taking advantage of hardware acceleration significantly closes the gap between pure emulation and native, but that different implementations and host graphics stacks show distinct variation. The microbenchmarks show that our architecture amplifies the overheads in the traditional graphics API bottlenecks: draw calls, downloading buffers, and batch sizes.

Our virtual GPU architecture runs modern graphics-intensive games and applications at interactive frame rates while preserving virtual machine portability. The applications we tested achieve from 86% to 12% of native rates and 43 to 18 frames per second with VMware Fusion 2.0.

1 Introduction

Over the past decade, virtual machines (VMs) have become increasingly popular as a technology for multiplexing both desktop and server commodity x86 computers. Over that time, several critical challenges in CPU virtualization were solved and there are now both software and hardware techniques for virtualizing CPUs with very low overheads [1]. I/O virtualization, however, is still very much an open problem and a wide variety of strategies are used. Graphics co-processors (GPUs) in particular present a challenging mixture of broad complexity, high performance, rapid change, and limited documentation.

Modern high-end GPUs have more transistors, draw more power, and offer at least an order of magnitude

more computational performance than CPUs. At the same time, GPU acceleration has extended beyond entertainment (e.g., games and video) into the basic windowing systems of recent operating systems and is starting to be applied to non-graphical high-performance applications including protein folding, financial modeling, and medical image processing. The rise in applications that exploit, or even assume, GPU acceleration makes it increasingly important to expose the physical graphics hardware in virtualized environments. Additionally, virtual desktop infrastructure (VDI) initiatives have led many enterprises to try to simplify their desktop management by delivering VMs to their users. Graphics virtualization is extremely important to a user whose primary desktop runs inside a VM.

GPUs pose a unique challenge in the field of virtualization. Machine virtualization multiplexes physical hardware by presenting each VM with a *virtual device* and combining their respective operations in the hypervisor platform in a way that utilizes native hardware while preserving the illusion that each guest has a complete stand-alone device. Graphics processors are extremely complicated devices. In addition, unlike CPUs, chipsets, and popular storage and network controllers, GPU designers are highly secretive about the specifications for their hardware. Finally, GPU architectures change dramatically across generations and their generational cycle is short compared to CPUs and other devices. Thus, it is nearly intractable to provide a virtual device corresponding to a real modern GPU. Even starting with a complete implementation, updating it for each new GPU generation would be prohibitively laborious. Thus, rather than modeling a complete modern GPU, our primary approach paravirtualizes: it delivers an idealized software-only GPU and our own custom graphics driver for interfacing with the guest operating system.

The main technical contributions of this paper are (1) a taxonomy of GPU virtualization strategies—both emulated and passthrough-based, (2) an overview of the virtual graphics stack in VMware's hosted architecture, and (3) an evaluation and comparison of VMware Fusion's 3D acceleration with other approaches. We find that a hosted model [2] is a good fit for handling complicated, rapidly changing GPUs while the largely asynchronous

graphics programming model is still able efficiently to utilize GPU hardware acceleration.

The rest of this paper is organized as follows. Section 2 provides background and some terminology. Section 3 describes a taxonomy of strategies for exposing GPU acceleration to VMs. Section 4 describes the device emulation and rendering thread of the graphics virtualization in VMware products. Section 5 evaluates the 3D acceleration in VMware Fusion. Section 6 summarizes our findings and describes potential future work.

2 Background

While CPU virtualization has a rich research and commercial history, graphics hardware virtualization is a relatively new area. VMware's virtual hardware has always included a display adapter, but it initially included only basic 2D support [3]. Experimental 3D support did not appear until VMware Workstation 5.0 (April 2005). Both Blink [4] and VMGL [5] used a user-level Chromium-like approach [6] to accelerate fixed function OpenGL in Linux and other UNIX-like guests. Parallels Desktop 3.0 [7] accelerates some OpenGL and Direct3D guest applications with a combination of Wine and proprietary code [8], but loses its interposition while those applications are running. Finally, at the most recent Intel Developer Forum, Parallels presented a demo that dedicates an entire native GPU to a single virtual machine using Intel's VT-d [9, 10].

The most immediate application for GPU virtualization is to desktop virtualization. While server workloads still form the core use case for virtualization, desktop virtualization is now the strongest growth market [11]. Desktop users run a diverse array of applications, including entertainment, CAD, and visualization software. Windows Vista, Mac OS X, and recent Linux distributions all include GPU-accelerated windowing systems. Furthermore, an increasing number of ubiquitous applications are adopting GPU acceleration. Adobe Flash Player 10, the next version of a product which currently reaches 99.0% of Internet viewers [12], will include GPU acceleration. There is a user expectation that virtualized applications will "just work", and this increasingly includes having access to their graphics card.

2.1 GPU Hardware

This section will briefly introduce GPU hardware. It is not within the scope of this paper to provide a full discussion of GPU architecture and programming models.

Graphics hardware has experienced a continual evolution from mere CRT controllers to powerful programmable stream processors. Early graphics accelerators could draw rectangles or bitmaps. Later graphics accelerators could rasterize triangles and transform and light them in hardware. With current PC graphics hardware, formerly fixed-function transformation and shading has become generally programmable. Graphics applications use high-level Application Programming Interfaces (APIs) to configure the pipeline, and provide *shader* programs which perform application specific per-vertex and per-pixel processing on the GPU [13].

Future GPUs are expected to continue providing increased programmability. Intel recently announced its Larrabee [14] architecture, a potentially disruptive technology which follows this trend to its extreme.

With the recent exception of many AMD GPUs, for which open documentation is now available [15], GPU hardware is proprietary. NVIDIA's hardware documentation, for example, is a closely guarded trade secret. Nearly all graphics applications interact with the GPU via a standardized API such as Microsoft's DirectX or the vendor-independent OpenGL standard.

3 GPU Virtualization Taxonomy

This section explores the GPU virtualization approaches we have considered at VMware. We use four primary criteria for judging them: performance, fidelity, multiplexing, and interposition. The former two emphasize minimizing the cost of virtualization: users desire native performance and full access to the native hardware features. The latter two emphasize the added value of virtualization: virtualization is fundamentally about enabling *many* virtual instances of *one* physical entity and then hopefully using that abstraction to deliver secure isolation, resource management, virtual machine portability, and many other features enabled by insulating the guest from physical hardware dependencies.

We observe that different use cases weight the criteria differently—for example a VDI deployment values high VM-to-GPU consolidation ratios (e.g., multiplexing) while a consumer running a VM to access a game or CAD application unavailable on his host values performance and likely fidelity. A tech support person maintaining a library of different configurations and an IT administrator running server VMs are both likely to value portability and secure isolation (interposition).

Since these criteria are often in opposition (e.g., performance at the expense of interposition), we describe several possible designs. Rather than give an exhaustive list, we describe points in the design space which highlight interesting trade-offs and capabilities. At a high level, we group them into two categories: front-end (application facing) and back-end (hardware facing).

3.1 Front-end Virtualization

Front-end virtualization introduces a virtualization boundary at a relatively high level in the stack, and runs the graphics driver in the host/hypervisor. This approach does not rely on any GPU vendor- or model-specific de-

tails. Access to the GPU is entirely mediated through the vendor provided APIs and drivers on the host while the guest only interacts with software. Current GPUs allow applications many independent "contexts" so multiplexing is easy. Interposition is not a given—unabstracted details of the GPU's capabilities may be exposed to the virtual machine for fidelity's sake—but it is straightforward to achieve if desired. However, there is a performance risk if too much abstraction occurs in pursuit of interposition.

Front-end techniques exist on a continuum between two extremes: *API remoting*, in which graphics API calls are blindly forwarded from the guest to the external graphics stack via remote procedure call, and *device emulation*, in which a virtual GPU is emulated and the emulation synthesizes host graphics operations in response to actions by the guest device drivers. These extremes have serious disadvantages that can be overcome by intermediate solutions. Pure API remoting is simple to implement, but completely sacrifices interposition and involves wrapping and forwarding an extremely broad collection of entry points. Pure emulation of a modern GPU delivers excellent interposition and implements a narrower interface, but a highly complicated and underdocumented one.

Our hosted GPU acceleration employs front-end virtualization and is described in Section 4. Parallels Desktop 3.0, Blink, and VMGL are other examples of front-end virtualization. Parallels appears to be closest to pure API remoting, as VM execution state cannot be saved to disk while OpenGL or Direct3D applications are running. VMGL uses Chromium to augment its remoting with OpenGL state tracking and Blink implements something similar. This allows them suspend-to-disk functionality and reduces the amount of data which needs to be copied across the virtualization boundary.

3.2 Back-end Virtualization

Back-end techniques run the graphics driver stack inside the virtual machine with the virtualization boundary between the stack and physical GPU hardware. These techniques have the potential for high performance and fidelity, but multiplexing and especially interposition can be serious challenges. Since a VM interacts directly with proprietary hardware resources, its execution state is bound to the specific GPU vendor and possibly the exact GPU model in use. However, exposure to the native GPU is excellent for fidelity: a guest can likely exploit the full range of hardware abilities.

The most obvious back-end virtualization technique is *fixed pass-through*: the permanent association of a virtual machine with full exclusive access to a physical GPU. Recent chipset features, such as Intel's VT-d, make fixed pass-through practical without requiring any

special knowledge of a GPU's programming interfaces. However, fixed pass-through is not a general solution. It completely forgoes any multiplexing and packing machines with one GPU per virtual machine (plus one for the host) is not feasible.

One extension of fixed pass-through is *mediated pass-through*. As mentioned, GPUs already support multiple independent contexts and mediated pass-through proposes dedicating just a context, or set of contexts, to a virtual machine rather than an entire GPU. This allows multiplexing, but incurs two additional costs: the GPU hardware must implement contexts in a way that they can be mapped to different virtual machines with low overheads and the host/hypervisor must have enough of a hardware driver to allocate and manage GPU contexts. Potentially third, if each context does not appear as a full (logical) device, the guest device drivers must be able to handle it.

Mediated pass-through is still missing any interposition features beyond (perhaps) basic isolation. A number of tactics using paravirtualization or standardization of a subset of hardware interfaces can potentially unlock these additional interposition features. Analogous techniques for networking hardware were presented at VMworld 2008 [16].

4   VMware's Virtual GPU

All of VMware's products include a virtual display adapter that supports VGA and basic high resolution 2D graphics modes. On VMware's hosted products, this adapter also provides accelerated GPU virtualization using a front-end virtualization strategy. To satisfy our design goals, we chose a flavor of front-end virtualization which provides good portability and performance, and which integrates well with existing operating system driver models. Our approach is most similar to the *device emulation* approach above, but it includes characteristics similar to those of *API remoting*. The in-guest driver and emulated device communicate asynchronously with VMware's Mouse-Keyboard-Screen (MKS) abstraction. The MKS runs as a separate thread and owns all of our access to the host GPU (and windowing system in general).

4.1   SVGA Device Emulation

Our virtual GPU takes the form of an emulated PCI device, the *VMware SVGA II* card. No physical instances of this card exist, but our virtual implementation acts like a physical graphics card in most respects. The architecture of our PCI device is outlined by Figure 1. Inside the VM, it interfaces with a device driver we supply for common guest operating systems. Currently only the Windows XP driver has 3D acceleration support. Outside the VM, a user-level device emulation process is re-

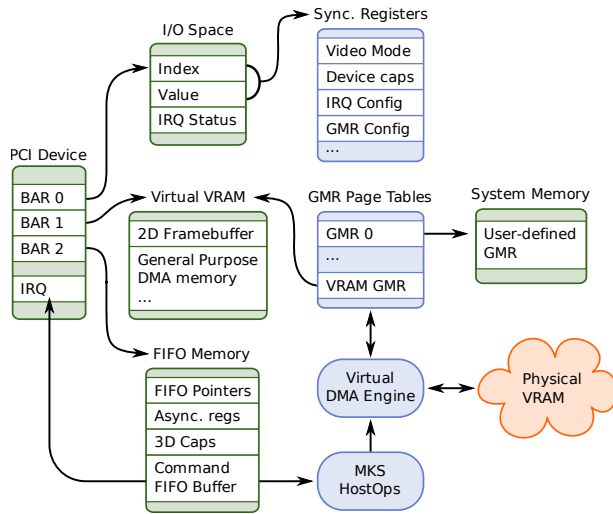sponsible for handling accesses to the PCI configuration and I/O space of the SVGA device.

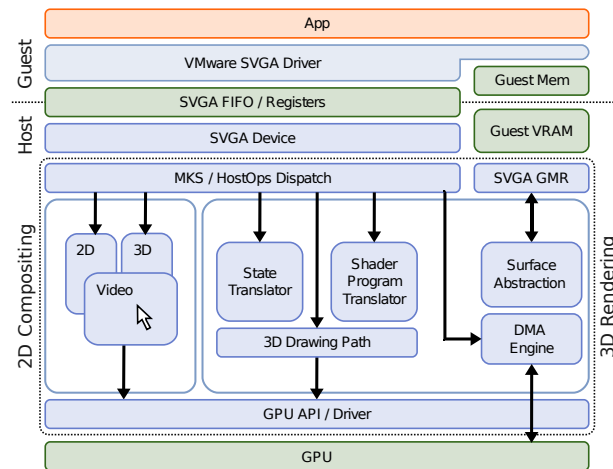Figure 1: *VMware SVGA II* device architecture

Figure 2: The virtual graphics stack. The MKS/HostOps Dispatch and rendering occur asynchronously in their own thread.

Our virtual graphics device provides three fundamental kinds of virtual hardware resources: registers, Guest Memory Regions (GMRs), and a FIFO command queue.

Registers may be located in I/O space, for infrequent operations that must be emulated synchronously, or in the faster "FIFO Memory" region, which is backed by plain system memory on the host. I/O space registers are used for mode switching, GMR setup, IRQ acknowledgement, versioning, and for legacy purposes. FIFO registers include large data structures, such as the host's 3D rendering capabilities, and fast-moving values such as the mouse cursor location—this is effectively a shared

memory region between the guest driver and the MKS.

GMRs are an abstraction for guest owned memory which the virtual GPU is allowed to read or write. GMRs can be defined by the guest's video driver using arbitrary discontiguous regions of guest system memory. Additionally, there always exists one default GMR: the device's "virtual VRAM." This VRAM is actually host system memory, up to 128 MB, mapped into PCI memory space via BAR1. The beginning of this region is reserved as a 2D framebuffer.

In our virtual GPU, physical VRAM is not directly visible to the guest. This is important for portability, and it is one of the primary trade-offs made by our frontend virtualization model. To access physical VRAM *surfaces* like textures, vertex buffers, and render targets, the guest driver schedules an asynchronous DMA operation which transfers data between a surface and a GMR. In every surface transfer, this DMA mechanism adds at least one copy beyond the normal overhead that would be experienced in a non-virtualized environment or with back-end virtualization. Often only this single copy is necessary, because the MKS can provide the host's OpenGL or Direct3D implementation with direct pointers into mapped GMR memory. This virtual DMA model has the potential to far outperform a pure API remoting approach like VMGL or Chromium, not only because so few copies are necessary, but because the guest driver may cache lockable Direct3D buffers directly in GMR memory.

Like a physical graphics accelerator, the SVGA device processes commands asynchronously via a lockless FIFO queue. This queue, several megabytes in size, occupies the bulk of the FIFO Memory region referenced by BAR2. During unaccelerated 2D rendering, FIFO commands are used to mark changed regions in the framebuffer, informing the MKS to copy from the guest framebuffer to the physical display. During 3D rendering, the FIFO acts as transport layer for our architecture-independent SVGA3D rendering protocol. FIFO commands also initiate all DMA operations, perform hardware accelerated blits, and control accelerated video and mouse cursor overlays.

We deliver host to guest notifications via a virtual interrupt. Our virtual GPU has multiple interrupt sources which may be programmed via FIFO registers. To measure the host's command execution progress, the guest may insert FIFO *fence* commands, each with a unique 32-bit ID. Upon executing a fence, the host stores its value in a FIFO register and optionally delivers an interrupt. This mechanism allows the guest to very efficiently check whether a specific command has completed yet, and to optionally wait for it by sleeping until a *FIFO goal* interrupt is received.

The SVGA3D protocol is a simplified and idealized

adaptation of the Direct3D API. It has a minimal number of distinct commands. Drawing operations are expressed using a single flexible vertex/index array notation. All host VRAM resources, including 2D textures, 3D textures, cube environment maps, render targets, and vertex/index buffers are represented using a homogeneous surface abstraction. Shaders are written in a variant of Direct3D's bytecode format, and most fixed-function render states are based on Direct3D render state.

This protocol acts as a common interchange format for GPU commands and state. The guest contains API implementations which produce SVGA3D commands rather than commands for a specific GPU. This provides an opportunity to actively trade capability for portability. The host can control which of the physical GPU's features are exposed to the guest. As a result, VMs using SVGA3D are widely portable between different physical GPUs. It is possible to suspend a live application to disk, move it to a different host with a different GPU or MKS backend, and resume it. Even if the destination GPU exposes fewer capabilities via SVGA3D, in some cases our architecture can use its layer of interposition as an opportunity to emulate missing features in software. This portability assurance is critical for preventing GPU virtualization from compromising the core value propositions of machine virtualization.

4.2 Rendering

This FIFO design is inherently asynchronous. All host-side rendering happens in the MKS thread, while the guest's virtual CPUs execute concurrently. As illustrated in Figure 2, access to the physical GPU is mediated first through the GPU vendor's driver running in the host OS, and secondly via the Host Operations (HostOps) backends in the MKS. The MKS has multiple HostOps backend implementations including GDI and X11 backends to support basic 2D graphics on all Windows and Linux hosts, a VNC server for remote display, and 3D accelerated backends written for both Direct3D and OpenGL. In theory we need only an OpenGL backend to support Windows, Linux, and Mac OS hosts; however we have found Direct3D drivers to be of generally better quality, so we use them when possible. Additional backends could be written to access GPU hardware directly.

The guest video driver writes commands into FIFO memory, and the MKS processes them continuously on a dedicated rendering thread. This design choice is critical for performance, however it introduces several new challenges in synchronization. In part, this is a classic producer-consumer problem. The FIFO requires no host-guest synchronization as long as it is never empty nor full, but the host must sleep any time the FIFO is empty, and the guest must sleep when it is full. The guest may also need to sleep for other reasons. The guest

video driver must implement some form of flow control, so that video latency is not unbounded if the guest submits FIFO commands faster than the host completes them. The driver may also need to wait for DMA completion, either to recycle DMA memory or to read back results from the GPU. To implement this synchronization efficiently, the FIFO requires both guest to host and host to guest notifications.

The MKS will normally poll the command FIFO at a fixed rate, between 25 and 100 Hz. This is effectively the virtual vertical refresh rate of the device during unaccelerated 2D rendering. During synchronization-intensive 3D rendering, we need a lower latency guest to host notification. The guest can write to the *doorbell*, a register in I/O space, to explicitly ask the host to poll the command FIFO immediately.

5 Evaluation

We conducted two categories of tests: application benchmarks, and microbenchmarks. All tests were conducted on the same physical machine: a 2^{nd} generation Apple Mac Pro, with a total of eight 2.8 GHz Intel Xeon cores and an ATI Radeon HD2600 graphics card. All VMs used a single virtual CPU. With one exception, we found that all non-virtualized tests were unaffected by the number of CPU cores enabled.

5.1 Application Benchmarks

Application	Resolution	FPS
RTHDRIBL	1280×1024	22
RTHDRIBL	640×480	27.5
Half Life 2: Episode 2	1600×1200	22.2
Half Life 2: Episode 2	1024×768	32.2
Civilization 4	1600×1200	18
Max Payne 2	1600×1200	42

Table 1: Absolute frame rates with VMware Fusion 2.0. All applications run at interactive speeds (18–42 FPS).

The purpose of graphics acceleration hardware is to provide higher performance than would be possible using software alone. Therefore, in this section we will measure both the performance impact of virtualized graphics relative to non-virtualized GPU hardware, and the amount of performance improvement relative to TransGaming's SwiftShader [17] software renderer, running in a VMware Fusion virtual machine.

In addition to VMware Fusion 2.0, which uses the architecture described above, we measured Parallels Desktop 3.0 where possible (three of our configurations do not run). Both versions are the most recent public release at time of writing. To demonstrate the effects that can be caused by API translation and by the host graphics stacks, we also ran our applications on VMware Workstation 6.5. These used our Direct3D rendering backend
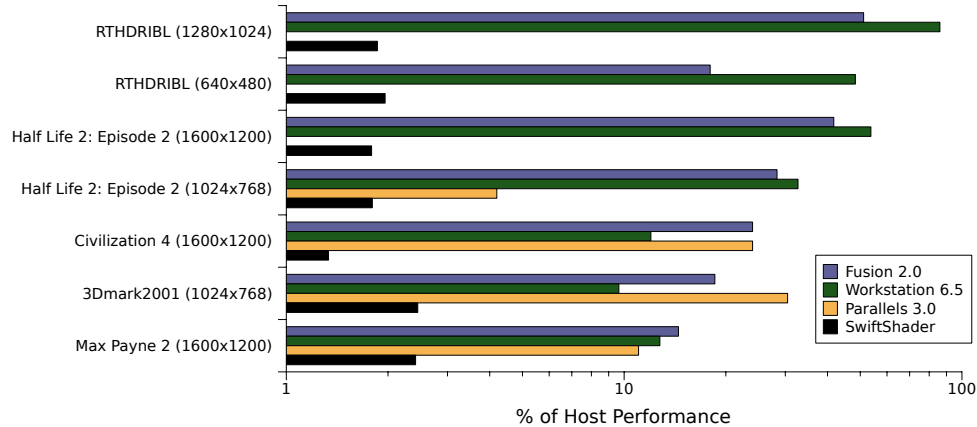
Figure 3: Relative performance of software rendering (SwiftShader) and three hardware accelerated virtualization techniques. The log scale highlights the huge gap between software and hardware acceleration versus the gap between virtualized and native hardware.

on the same hardware, but running Windows XP using Boot Camp.

It is quite challenging to measure the performance of graphics virtualization implementations accurately and fairly. The system under test has many variables, and they are often difficult or impossible to isolate. The virtualized operating system, host operating system, CPU virtualization overhead, GPU hardware, GPU drivers, and application under test may each have a profound effect on the overall system performance. Any one of these components may have opaque fast and slow paths—small differences in the application under test may cause wide gaps in performance, due to subtle and often hidden details of each component's implementation. For example, each physical GPU driver may have different undocumented criteria for transferring vertex data at maximum speed.

Additionally, the matrix of possible tests is limited by incompatible graphics APIs. Most applications and benchmarks are written for a single API, either OpenGL or Direct3D. Each available GPU virtualization implementation has a different level of API support. Parallels Desktop supports both OpenGL and Direct3D, VMware Fusion supports only Direct3D applications, and VMGL supports only OpenGL.

Figure 3 summarizes the application benchmark results. All three virtualization products performed substantially better than the fastest available software renderer, which obtained less than 3% of native performance in all tests. Applications which are mostly GPU limited, RTHDRIBL [18] and *Half Life 2: Episode 2*, ran at closer to native speeds. Max Payne exhibits low performance relative to native, but that reflects the low ratio of GPU load to API calls. As a result, virtualization overhead occupies a higher proportion of the whole

execution time. In absolute terms, though, Max Payne has the highest frame rate of our applications.

Table 1 reports the actual frame rates exhibited with these applications under VMware Fusion. While our virtualized 3D acceleration still lags native performance, we make two observations: it still achieves interactive frame rates and it closes the lion's share of the gulf between software rendering and native performance. For example, at 1600×1200, VMware Fusion renders Half-Life 2 at 22 frames per second, which is 23.35x faster than software rendering and only 2.4x slower than native.

5.2 Microbenchmarks

To better understand the nature of front-end virtualization's performance impact, we performed a suite of microbenchmarks based on triangle rendering speed under various conditions. For all microbenchmarks, we rendered unlit untextured triangles using Vertex Shader 1.1 and the fixed-function pixel pipeline. This minimizes our dependency on shader translation and GPU driver implementation.

Each test renders a fixed number of frames, each containing a variable number of draw calls with a variable length vertex buffer. For security against jitter or drift caused by timer virtualization, all tests measured elapsed time via a TCP/IP server running on an idle physical machine. Parameters for each test were chosen to optimize frame duration, so as to minimize the effects of noise from time quantization, network latency, and vertical sync latency.

The static vertex test, Figure 4(a), tests performance scalability when rendering vertex buffers which do not change contents. In Direct3D terms, this tests the managed buffer pool. Very little data must be exchanged

between host and guest in this test, so an ideal front-end virtualization implementation would do quite well. VMware Workstation manages to get just over 80% of the host's performance in this test. Parallels Desktop and VMware Fusion get around 30%. In our experience, this is due to inefficiency in the Vertex Buffer Object support within Mac OS's OpenGL stack.

The dynamic vertex test, Figure 4(b), switches from the managed buffer pool back to the default Direct3D buffer pool, and uploads new vertex data prior to each of the 100 draws per frame. It tests the driver stack's ability to stream data to the GPU, and manage the re-use of buffer memory.

The next test, Figure 4(c), is intended to test virtualization overhead while performing a GPU-intensive operation. While triangles in previous tests had zero pixel coverage, this tests renders triangles covering half the viewport. Ideally, this test would show nearly identical results for any front-end virtualization implementation. The actual results are relatively close, but on VMware's platform there is a substantial amount of noise in the results. This appears to be due to the irregular completion of asynchronous commands when the physical GPU is under heavy load. Also worth noting is the fact that VMware Fusion, on average, performed better than the host machine. It's possible that this test is exercising a particular drawing state which is more optimized in ATI's Mac OS OpenGL driver than in their Windows Direct3D driver.

The final test, Figure 4(d), measures the overhead added to every separate draw. This was the only test where we saw variation in host performance based on the number of enabled CPU cores. This microbenchmark illustrates why the number of draw calls per frame is, in our experience, a relatively good predictor of overall application performance with front-end GPU virtualization.

6   Conclusion

In VMware's hosted architecture, we have implemented front-end GPU virtualization using a virtual device model with a high level rendering protocol. We have shown it to run modern graphics-intensive games and applications at interactive frame rates while preserving virtual machine interposition.

There is much future work in developing reliable benchmarks which specifically stress the performance weaknesses of a virtualization layer. Our tests show API overheads of about 2 to 120 times that of a native GPU. As a result, the performance of a virtualized GPU can be highly dependent on subtle implementation details of the application under test.

Back-end virtualization holds much promise for performance, breadth of GPU feature support, and ease of driver maintenance. While *fixed pass-through* is easy, none of the more advanced techniques have been demonstrated. This is a substantial opportunity for work by GPU and virtualization vendors.
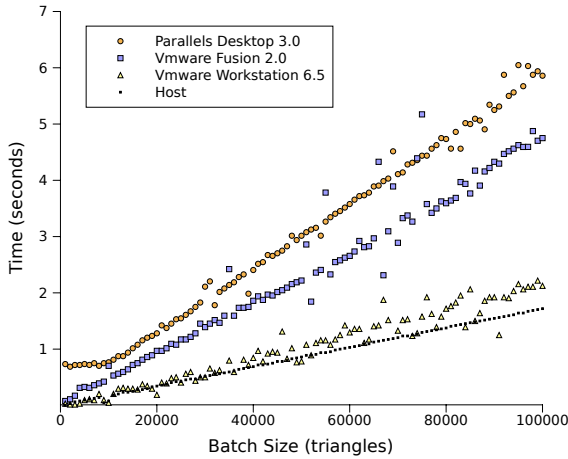
Front-end virtualization currently shows a substantial degradation in performance and GPU feature set relative to native hardware. Nevertheless, it is already enabling virtualized applications to run interactively that could never have been virtualized before, and is a foundation for virtualization of tomorrow's GPU-intensive software. Even as back-end virtualization gains popularity, front-end virtualization can fill an important role for VMs which must be portable among diverse GPUs.
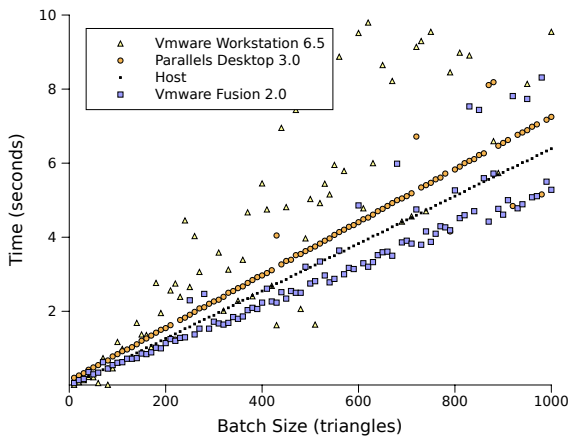
7   Acknowledgements

References

[1] Keith Adams and Ole Agesen. A Comparison of Software and Hardware Techniques for x86 Virtualization. In *Proceedings of ASPLOS '06*, October 2006.

[2] Jeremy Sugerman, Ganesh Venkitachalam, and Beng-Hong Lim. Virtualizing I/O Devices on VMware Workstation's Hosted Virtual Machine Monitor. In *Proceedings of the 2001 USENIX Annual Technical Conference*, June 2001.

[3] VMware SVGA Device Interface and Programming Model. In X.org source repository, `xf86-video-vmware` driver README.

[4] Jacob Gorm Hansen, Blink: Advanced Display Multiplexing for Virtualized Applications. In *Proceedings of NOSSDAV 2007*, June 2007.

[5] H. Andrés Lagar-Cavilla et al., VMM-Independent Graphics Acceleration. In *Proceedings of VEE '07*.

[6] Greg Humphreys et al., Chromium: A Stream-Processing Framework for Interactive Rendering on Clusters. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, pages 693-702 (2002).

[7] Parallels Desktop, `http://www.parallels.com/en/desktop/`

[8] Parallels on the Wine project wiki, `http://wiki.winehq.org/Parallels`
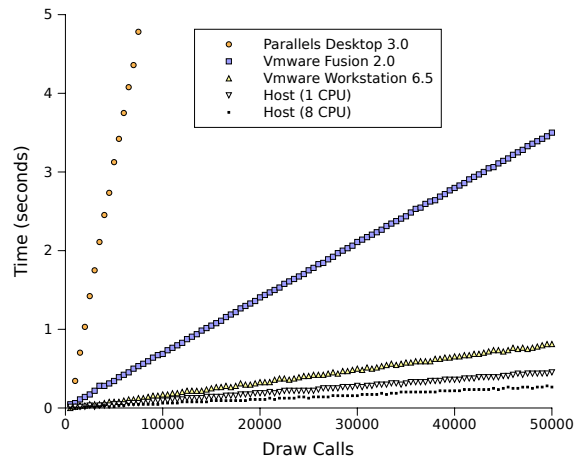
(a) Static vertex rendering performance

(b) Dynamic vertex rendering performance

(c) Filled triangle rendering performance

(d) Draw call overhead

Figure 4: Microbenchmark results

[9] IDF SF08: Parallels and Intel Virtualization for Directed I/O, http://www.youtube.com/watch?v=EiqMR5Wx_r4

[10] D. Abramson et al., Intel Virtualization Technology for Directed I/O. In *Intel Technology Journal*, http://www.intel.com/technology/itj/2006/v10i3/ (August 2006).

[11] Andi Mann, *Virtualization and Management: Trends, Forecasts, and Recommendations*, Enterprise Management Associates. (2008).

[12] Flash Player Penetration, http://www.adobe.com/products/player_census/flashplayer/

[13] John Owens, GPU architecture overview. In *SIGGRAPH '07: ACM SIGGRAPH 2007 courses*, http://doi.acm.org/10.1145/1281500.1281643

[14] Larry Seiler et al., Larrabee: A Many-Core x86 Architecture for Visual Computing. In *ACM Transactions on Graphics* Vol. 27, No. 3, Article 18 (August 2008).

[15] AMD Developer Guides and Manuals, http://developer.amd.com/documentation/guides/Pages/default.aspx

[16] Howie Xu et al., TA2644: Networking I/O Virtualization, VMworld 2008.

[17] SwiftShader, http://www.transgaming.com/products/swiftshader/

[18] Real-Time High Dynamic Range Image-Based Lighting demo, http://www.daionet.gr.jp/~masa/rthdribl/