

VMGL Intro.

國網中心-格網技術組

Rock(郭文傑), Rider(涂哲源)

rock@nchc.org.tw, rider@nchc.org.tw

Why Virtualize 3D ?

- Virtualization

- Cost
- Green
- Convenience

- 3D

- Interaction
- Reality
- Recreation
- User-Interface



GPU Virtualization is difficult ?

- Closed information

- HW specs
- Device driver

- Lack of standard

- uncertain interface – AGP ? PCI ? PCI-X ? PCI-E ?
- unifying HW abstraction
- common standard – x86, IDE, SCSI ...

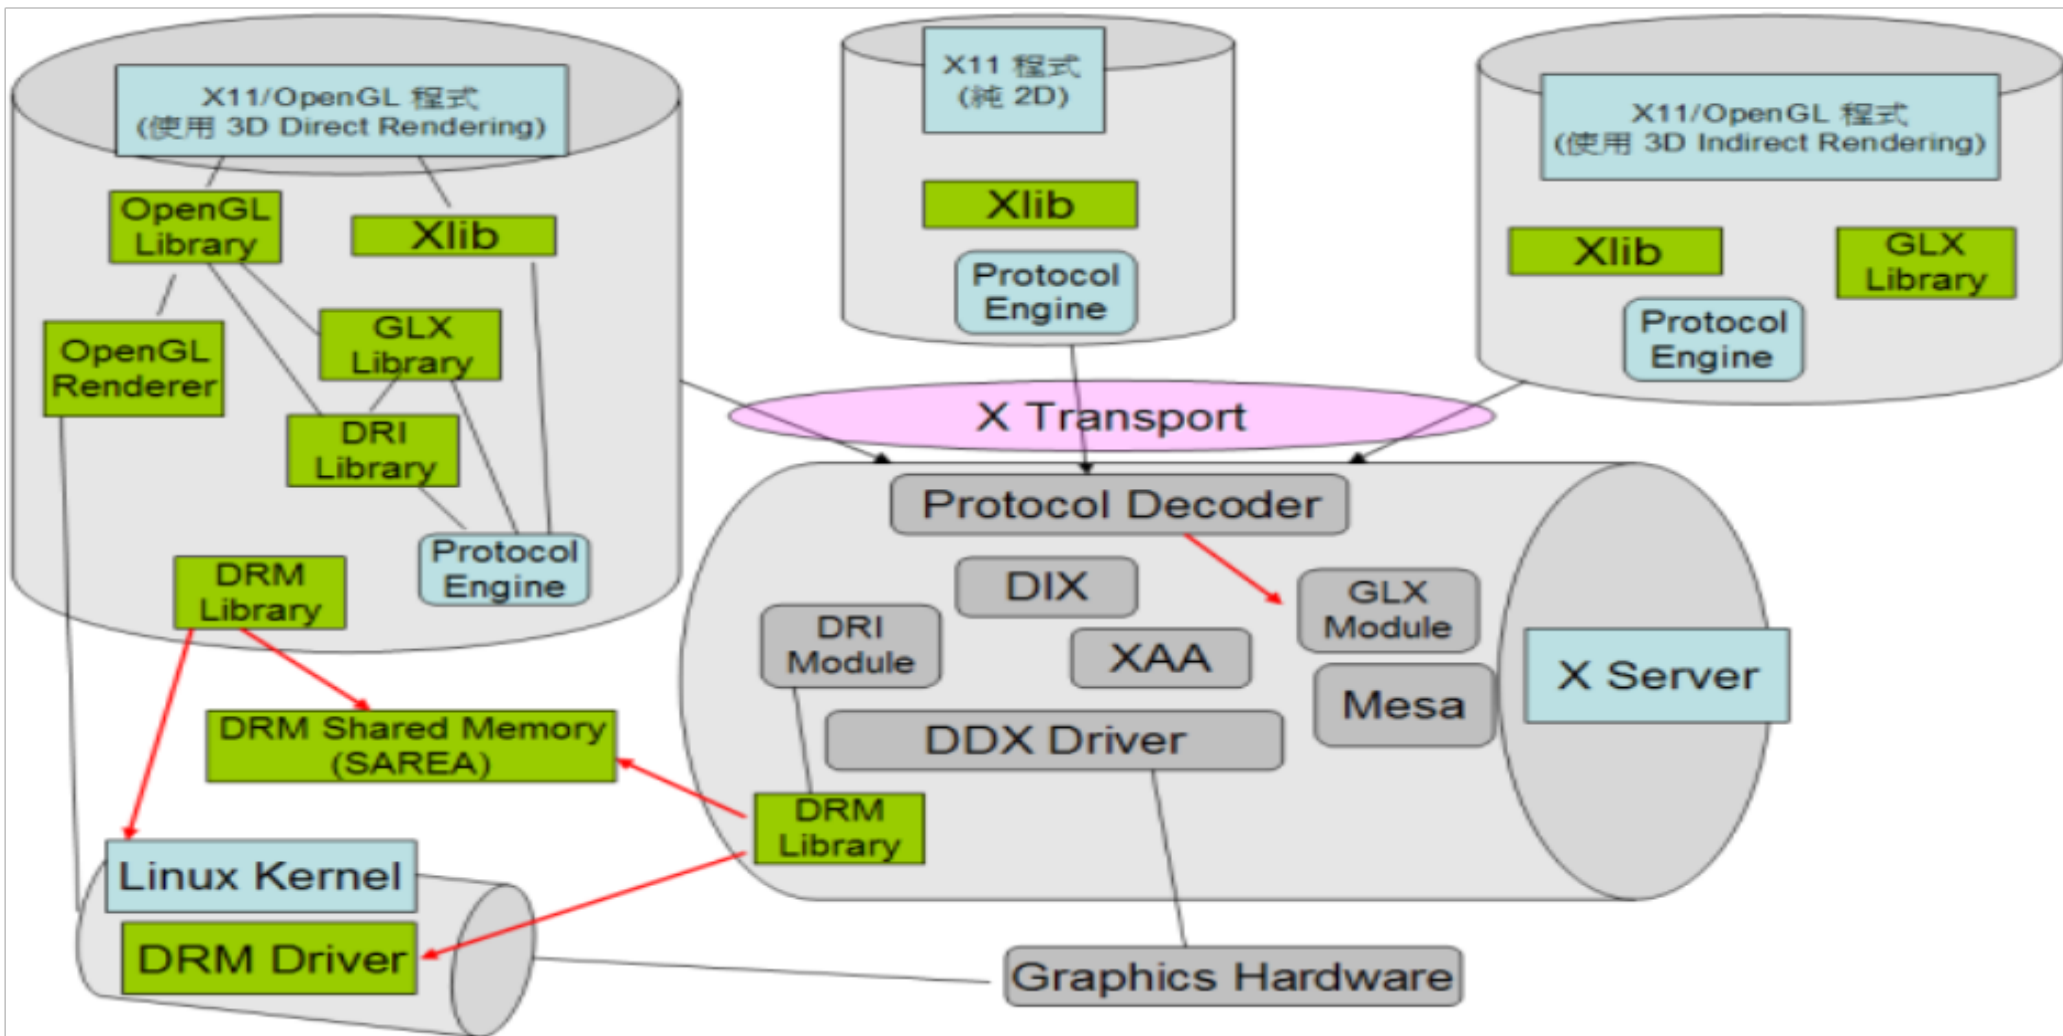
High-level API ?

3D Graphics API

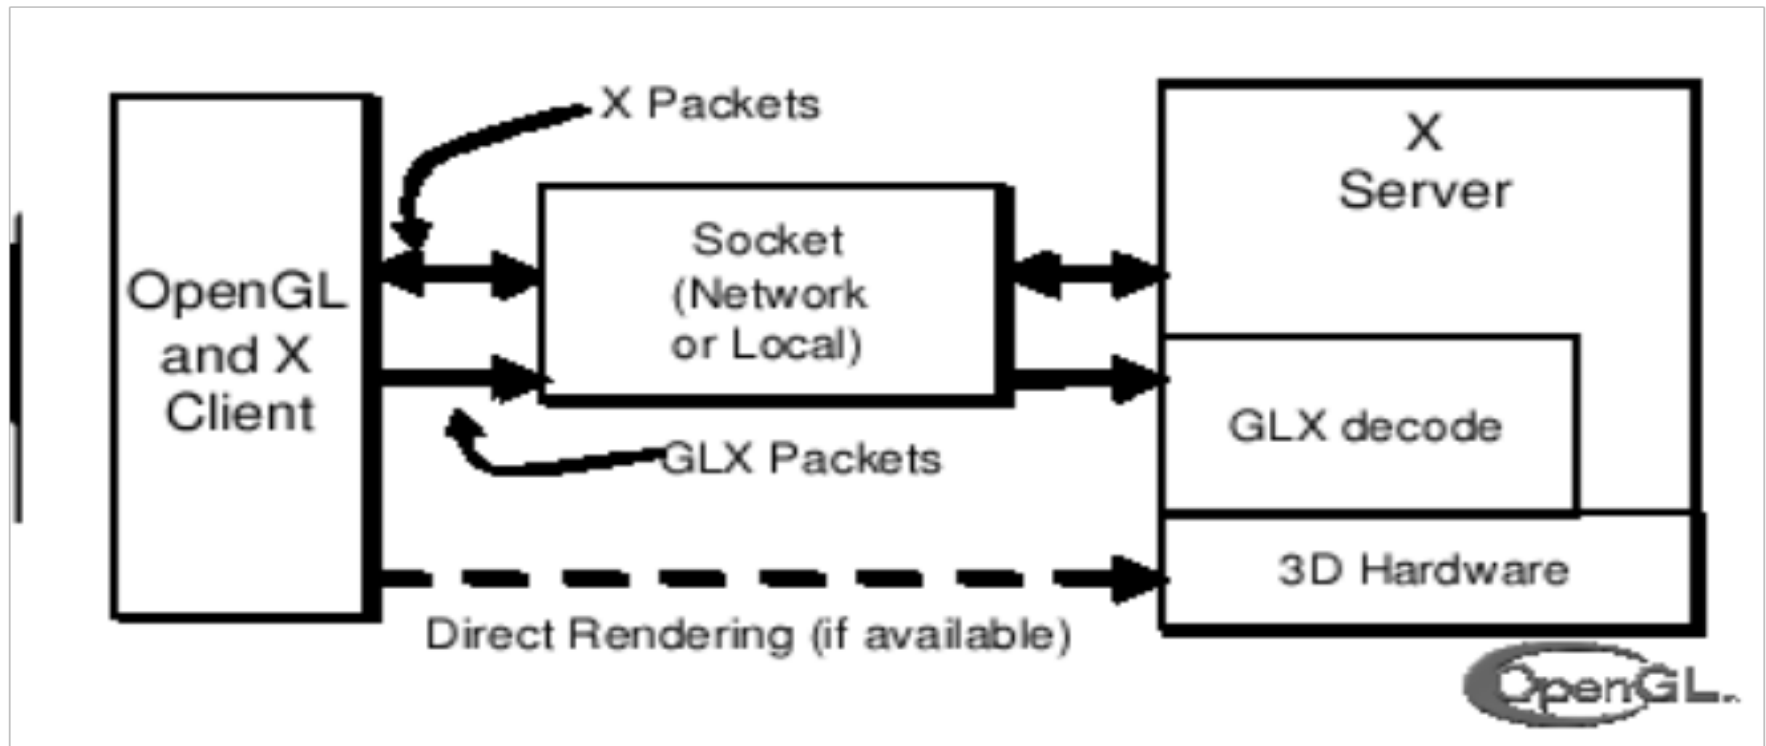
- Direct3D
- OpenGL



X11 / OpenGL 關係圖



OpenGL extension – GLX



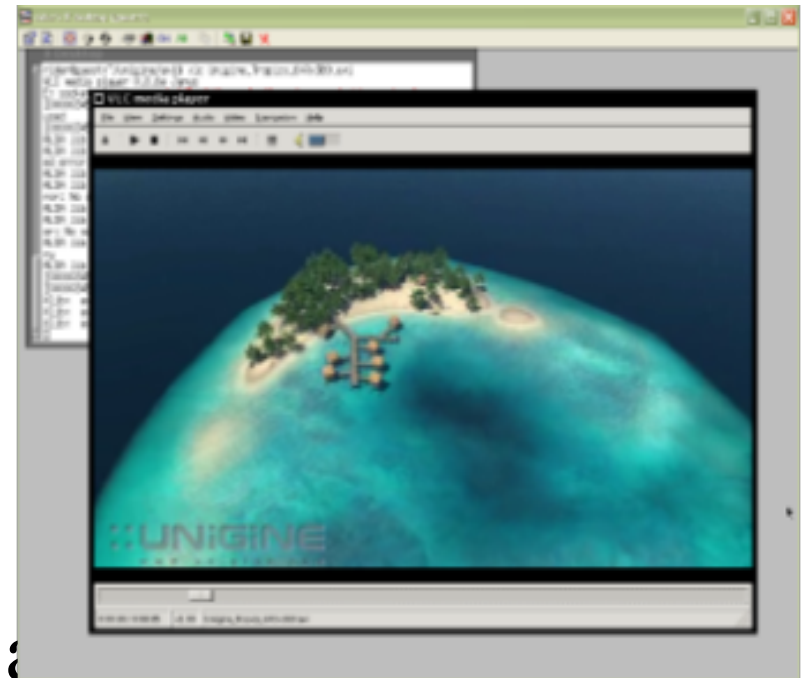
Requirement ?

- Graphics-intensive applications
- Excellent rendering performance
- Cross-Platform – VMs & GPUs
- Suspend & Resume across different GPUs
- Lightweight User-Interface for thin client



VMGL: Virtualized OpenGL

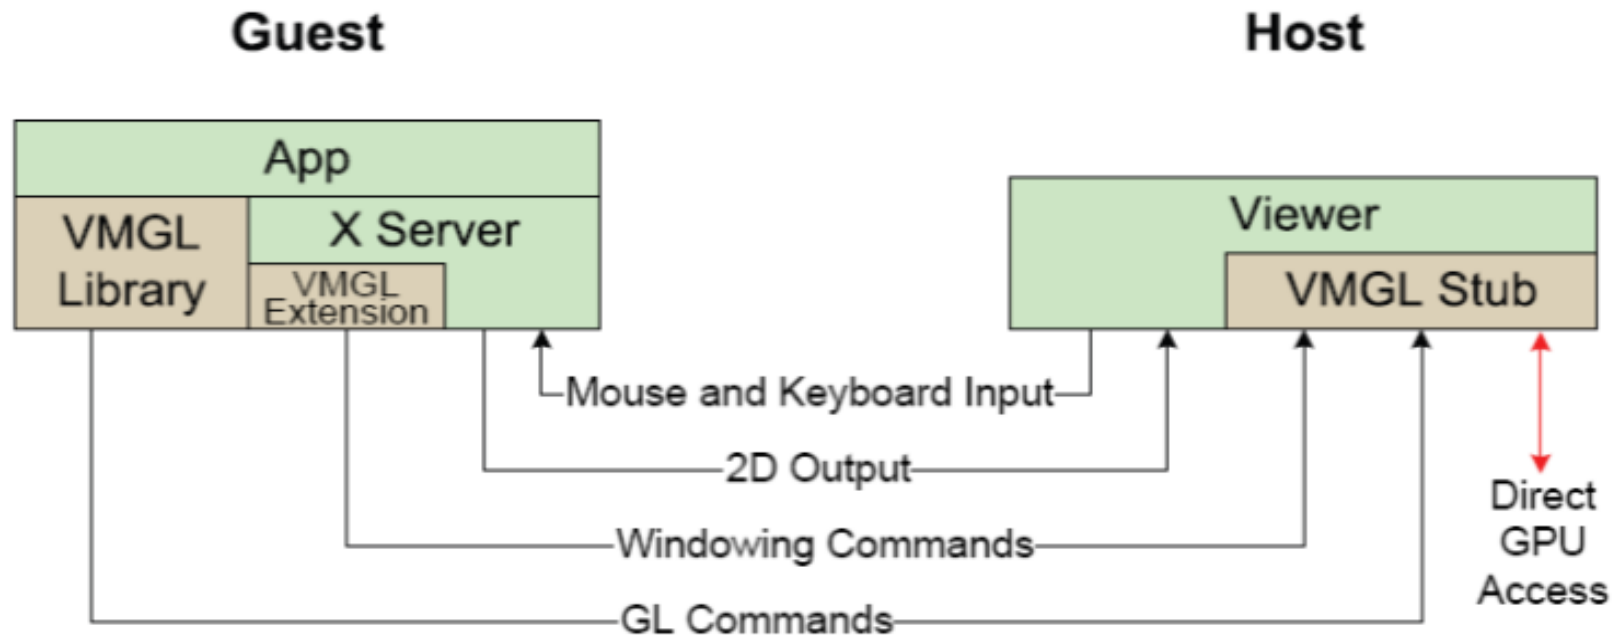
- OpenGL virtualization
- HW accelerated rendering
- VM Host
 - vendor-supplied GPU driver
 - OpenGL library
- Cross-Platform
- 87% or better of native HW &
- Suspend & Resume




Architecture features - VMGL

- Virtualizing the OpenGL API
 - Portability
 - Compatibility
- Use a Network Transport
 - Cross-VM

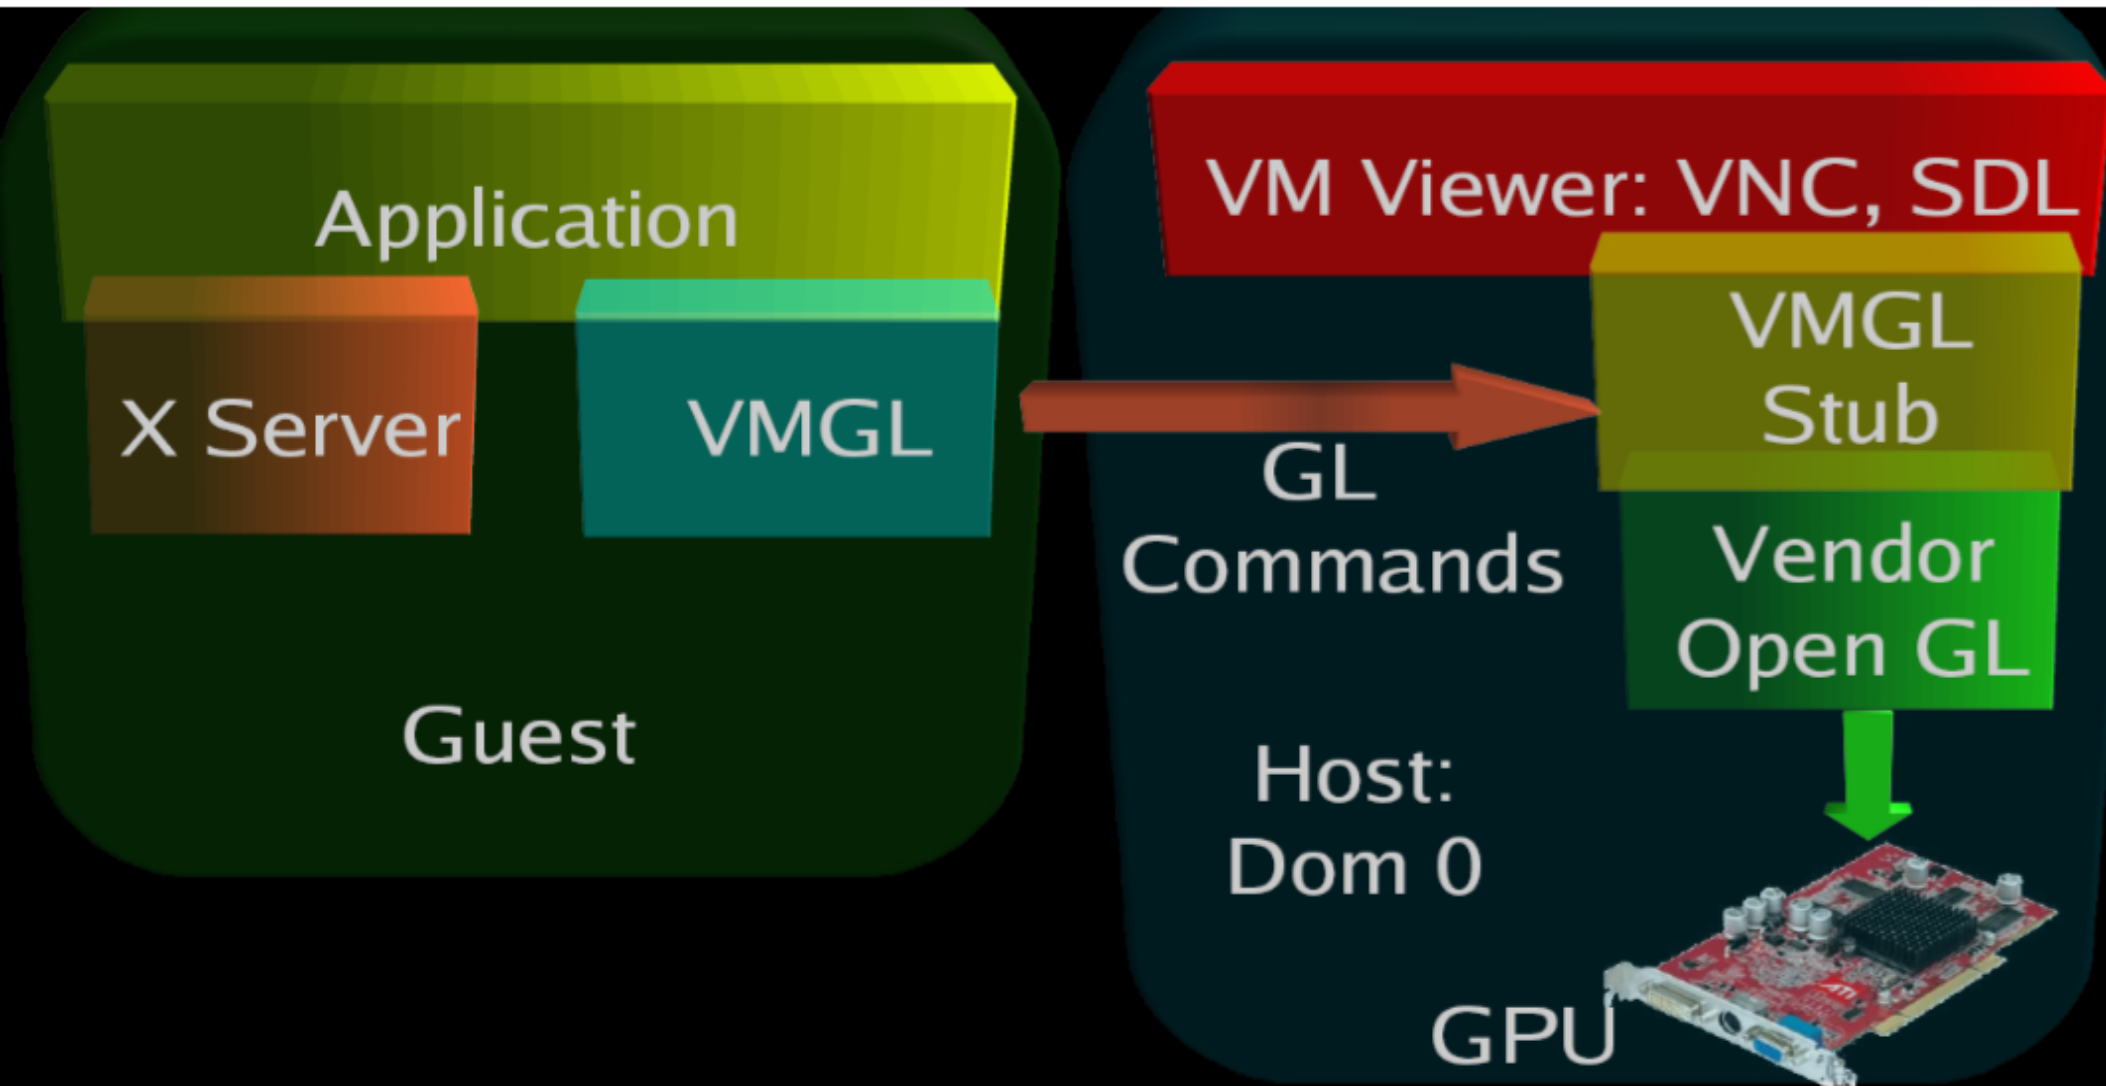
VMGL Architecture



OpenGL for X11 apps



VMGL for X11 apps



VMGL Suspend & Resume

- Shadow driver
 - Keep OpenGL state
- OpenGL contexts
 - Global Context State
 - Texture State
 - Display Lists

Dom 0 GPU Drivers

- **ATI & NVIDIA:**
 - GPU Mem mapping in user-space GL lib
- **Oblivious to Xen additional indirection**
 - Virtual \rightarrow Physical(VM) \rightarrow Machine
 - Even for domain 0
- **Fix open source portion of driver**
- **Use Xen-paravirt mem mapping functions**

VMGL Evaluation

- **VMGL: OpenGL Virtualization**
 - API v1.5
- Frames per second
- CPU, bandwidth consumption
- Resume latency, state size

Conclusion

- Strengths
 - Cross-Platform(VMs), Lightweight GUI Design
- Weaknesses
 - Network transparent, GL extensions
- Applications
 - Thin desktop, WebOS
- New trends
 - Intel VT-d (HW), NVIDIA SLI Multi-OS

Reference

● NVIDIA

- http://www.nvidia.com/object/cuda_home.html
- <http://forums.nvidia.com/index.php?>
- David Kirk/NVIDIA and Wen-mei W. Hwu, 2007 ECE 498AL, University of Illinois, Urbana-Champaign

● GViM

- <http://www.cc.gatech.edu/~vishakha/files/GViM.pdf>

● VMGL

- <http://www.cs.toronto.edu/~andreslc/xen-gl/>
- <http://www.cs.toronto.edu/~andreslc/publications/LagarCavillaVEE07.pdf>
- <http://www.cs.toronto.edu/~andreslc/publications/slides/Xen-Summit-2007/vmgl.pdf>

● Xen

- <http://trac.nchc.org.tw/grid/wiki/Reading/XenP>
- http://www.virtuatopia.com/index.php/Xen_Virtualization_Essentials

Discussion

