Proceedings of 2014 International Symposium on Embedded Technology (ISET 2014)
Seocho-gu Seoul, Korea, 22-23 May, 2014.

The GPGPU virtualization performance in actual network environment

Geun-yeong Bak* ShinHyoung Lee† Chuck Yoo†

* Graduate School of Convergence IT
Korea University, Seoul, Republic of Korea
E-mail: gypark@os.korea.ac.kr

† Collage of Information and Communication
Korea University, Seoul, Republic of Korea
E-mail: { shlee, hxy }@os.korea.ac.kr

Abstract

Due to development of GPU, GPGPU attempts to use the GPU device for high-performance computing rise. When apply the GPGPU attempts to virtualization environment, the inter-VM communication is inevitable because of nature of the GPU virtualization technology that the virtualized GPU can be used independently in only one VM. The inter-VM communication is significant overhead in virtualizing GPU. In this paper, we verified how the trade-off between improving performance achieved by GPU virtualization and the inevitable inter-VM overhead presents, implementing the GPGPU network routing structure of the relevant research in the actual network communication environment.

Keywords: savanna, GPU, GPGPU, Xen.

1. Introduction

Thanks to support for floating-point operations and the technology-intensive development, GPU becomes not only simple graphics processing unit but operational accelerating device. These changes in the GPU raise attempts to use the GPU device for high-performance computing. GPGPU[1] is a collective term for these attempts.

There are the tendencies to apply GPGPU system performance improvement to virtualized environments as well as general environments, and GPGPU performance improvement in the virtualized system can be realized by way of GPU virtualization. For GPU virtualization, the intel's VT-d hardware virtualization technology [2] emulates the GPU hardware to make VM (Virtual Machine) use the GPU device directly and exclusively.

Due to the nature of the VT-d technology that the virtualized GPU can be used independently in only one VM, a number of inter-VM communication is inevitable for multiple VM to share the GPU. That's because the virtual machine with virtualized GPU (DomG) must process all of the GPU-required tasks instead of other virtual machines without authorization to access the virtualized GPU (DomU). As a result of the method, sharing GPU with multiple VM could add the inter-VM communication overhead, and this overhead could act as a variable in deciding the system performance. Thus, if the GPU is used by multiple VMs, we should check the trade-off between VM performance improvements and overhead of inter-VM communication well.

In this paper, we verified how the trade-off presents implementing the GPGPU network routing structure of the relevant research in the actual network communication environment.

2. Relevant Researches

Implementing module that route virtual packets using virtualized GPU and return the results to client VM, the paper [3] shows measuring result of a trade-off between improved performance with GPGPU in a virtualized environment and inter-VM communication overhead: In the case that GPGPU handles the routing process, the routing rate is six times faster than in the otherwise case, but overhead occurring in the process of data transmission among VMs exceeds the amount of improved performance, so the overall performance declined greatly.

The paper [4], the subsequent paper of [3], suggests a structure for reducing inter-VM communication overhead (Savanna). With Savanna which was suggested in the paper [4], the whole time of GPGPU routing process including inter-VM communication could become lower than the general routing time.

3. Experiment

In this paper, with experiments using the forwarding module attached to Savanna, we verified how the system performance appears in the actual environment in which packets are transmitted among a network. The combined module (Savanna + original packet forwarding module) was operated in the same system as the one used in the paper [4], and the network communication environment is as follows.

First, the packets forwarding through the combined module had 13075000 pps (packet per second) rate when size of one packet was 64byte, and 820000 pps rate when one packet size was 1500byte (the maximum pps the packet sending node can pour to other nodes). All incoming packets received by the experiment module were sent to one node, and we used receiving pps rate of the node in which all packets come as a measure.The packet forwarding module used to experiments batched the 448 packets, the number of packets the GPU can treat at a time in parallel.

To see how the savanna affects in the performance in the actual network communication environment, the original packet forwarding module (origin. Module) and combined module were compared each other. The result of the comparison is presented in Table 1.

Table 1. Experiment Results

Packet size	Module	Send rate	Receive rate
64byte packet	Origin. Module	13075000 pps	4901809 pps
	Combined Module	13075000 pps	29702 pps
1500byte packet	Origin. Module	820000 pps	565603 pps
	Combined Module	820000 pps	65631 pps

4. Conclusion

As shown in the experimental results, savanna structure application to a real network environment makes significant performance degradation. Considering that the original forwarding module attached to Savanna is a forwarding-only version, the degradation of performance seems to be significant. Then, there are needs to analysis more detailed factors in the performance degradation and to design the ways to reduce degradation in performance.

5. References

[1] GPGPU, http://en.wikipedia.org/wiki/General-purpose_computing_on_graphics_processing_units

[2] VT-d, https://software.intel.com/en-us/articles/intel-virtualization-technology-for-directed-io-vt-d-enhancing-intel-platforms-for-efficient-virtualization-of-io-devices

[3] Yoo-kyung Uh, Chi-Young Lee, Chuck Yoo. 2013. "Communication Performance Evaluation in GPU Virtualized Environments." Korea Computer Congress, 63-65.

[4] Yoo-kyung Uh, Chuck Yoo, 2014. “Savanna: Xen 기반의 최소 통신을 이용한 GPGPU 가상화 기법”. Korea university master's thesis. http://www.riss.kr/link?id=T13383400