

## Dynamic memory mapping delivers additional flexibility to virtual resource management

WANG XiaoLin & LUO YingWei

Department of Computer Science and Technology, Peking University, Beijing 100871, China

The Department of Computer Science and Technology, Peking University, Beijing, China, has shown that a novel Dynamic Memory Mapping (DMM) model brings about additional flexibility to virtual resource management, leading to the feature-adjustable design of a virtual machine monitor (VMM). The study is reported in Issue 53 (June, 2010) of *SCIENCE CHINA Information Sciences* because of its significant research value.

Memory is one of the most frequently accessed components in virtual machine (VM) systems. Because a VM's memory requirement varies according to the running applications, disregarding the dynamic changes can result in suboptimal use of memory resources, which negatively affects the VM's performance. However, because the infrastructures of current technologies are usually independent of each other, they exhibit poor extensibility, integrity, and maintainability. To improve the flexibility and extensibility of the VMM, we need to implement a dynamic memory management mechanism in the VMM, while preserving the high efficiency of memory accesses from virtual machines.

To resolve these problems, this work proposes a DMM model [1]. The DMM model is a low-level memory management mechanism, which allows dynamic change of the mapping between the pseudo-physical memory as seen from VMs and the machine memory, while the virtual machine is running. On one hand, DMM is independent of, yet compatible with, various virtualization architectures, while on the other, it presents a uniform upward interface for supporting high-level memory management policies. As a result, the DMM layer incorporates high-level policies and low-level implementations by making both of them adjustable.

In this work, Prof. Wang, Prof. Luo and their group present the principle of the DMM model, and explain the procedures of various memory management policies under this model, such as demand paging, virtual memory and memory sharing. They also implement the DMM model in KVM, an open source VMM. They first designed a memory pool, a set of machine pages provided by the VMM to a particular virtual machine whose size could be expanded or shrunk while operating. To make the model work in a real system, they manipulated a page-level protection mechanism to propagate memory-mapping updates to the shadow page tables, which is the only way for a VM to access its virtualized memory. They also utilized reverse mapping, a data structure that maps a machine page back to all the shadow page table entries that have mapped it, to facilitate the mapping propagation.

The DMM model has several advantages over the current memory management mechanism in VMMs. The first is platform independence—the model is defined abstractly, thus is independent of implementations and computer architectures. The second is flexibility—the DMM model provides a uniform interface for integrating advanced memory management policies. Through the general mechanism, they can work together without conflict. Last but not least, the modular and layered design of the DMM reduces the complexity of a VMM's code base, and therefore improves the security and dependability of the system.

A journal reviewer noted: "This paper addresses the inefficiency in the design of current virtual machine monitors. Their approach is novel and systematic, and incurs only minor overheads. The result is of academic significance and practical value". Another reviewer said, "It enriches and expands the capacity and capability of virtualization. It offers us new methods to deploy and manage large numbers of virtual machines". A series of papers about virtual machine system optimization written by Prof. Wang, Prof. Luo and their group have been published in *SCIENCE CHINA Information Sciences* [2], *IEEE Cluster* [3], *ACM SIGOPS Operating System Review* [4].

The authors are affiliated to the Institute of Network Computing and Information Systems (NCIS, <http://ncis.pku.edu.cn>) at Peking University. This institute, led by Prof. LI XiaoMing, conducts research mainly in the fields of high productivity computing, search engine and Web mining, distributed systems, internet and mobile computing, and database technology.

*This research was supported by the National key Basic Research Program of China (2007CB310900), the National Natural Science Foundation of China (90718028 and 60873052), the National High Technology Research and Development Program of China (2008AA01Z112), and the MOE-Intel Information Technology Foundation (MOE-INTEL-08-09).*

- See the articles:** [1] Chen H G, Wang X L, Wang Z L, et al. DMM: A dynamic memory mapping model for virtual machines. *Sci China Inf Sci*, 2010, 53: 1097–1108, doi: 10.1007/s11432-010-3113-y  
[2] Wang X L, Sun Y F, Luo Y W, et al. Dynamic memory paravirtualization transparent to guest OS. *Sci China Inf Sci*, 2010, 53: 77–88, doi: 10.1007/s11432-010-0008-x  
[3] Luo Y W. Live and incremental whole system migration of virtual machines using block-bitmap. In: 2008 IEEE International Conference on Cluster Computing (Cluster'08), Tsukuba, Japan, 2008.9. 99–106  
[4] Zhao W M, Wang Z L, Luo Y W. Dynamic Memory Balancing for Virtual Machines. *ACM SIGOPS Operating System Reviews*, 2009, 43: 37–47