Proceedings of 2013 International Symposium on Embedded Technology(Interactive Session Only)
(ISET 2013)
Daegu Metropolitan City, 23-24 May, 2013.

A Compressed Page Sharing Between Virtual Machines

Jong-Hee Yun* Chi-Young Lee† Chuck Yoo†

*Department of Computer Science and Engineering
Korea University, Seoul, Republic of Korea
E-mail: yun.jonghee@gmail.com

† Department of Computer Science and Engineering
Korea University, Seoul, Republic of Korea
E-mail: {cylee, chuckyoo}@os.korea.ac.kr

Abstract

Nowadays, the virtualization system is being applied to embedded devices. However, embedded devices have limited memory. Therefore, we propose effective method in order to provide more memory for applications on the virtual machine. To realize purposed method, we implement a swapping method using a main memory and compression. When swapping occurs, we compress the swap page and store compressed swap page to another virtual machine with sufficient memory. If the compressed swap page already exists, it will be shared between virtual machines.

Keywords: embedded, virtualization, compressing swap page, memory sharing, xen.

1. Introduction

Main memory is the most important resource in the virtualization system. If a virtualized system runs many virtual machines, it causes a lack of memory. Thus, the virtual machines start swapping in order to make extra memory for an application. However, the swapping makes an overhead of disk I/O which decreases the performance of the application. In addition, most of embedded systems have the flash memory instead of a disk. However, it is difficult to use the flash memory for swapping. Firstly, a re-write operation on the flash memory is slow because it has to erase a block included a page to modify data of the page. Secondly, frequent swapping shortens the life time of flash memory. Therefore, a new swapping method for the virtual machines is required in order to overcome disadvantages of disk and flash memory.

2. A compressed swap page sharing Technique

Previous solution compresses swap pages and stores them into their memory [1]. This solution can reduce overhead of disk I/O; however, a lack of memory issue is still remained because of memory consumption by compressed pages.

Our method stores compressed swap pages into the memory on other virtual machine that has more enough memory. In addition, we share the compressed swap pages between many virtual machines.

We implement the method on Xen based virtualization system [2]. Xen has two types of domains. One domain (called domainU) is a virtual machine that can run applications. Another domain (called domain0) is responsible to deal with requests for hardware resource. We use the domain0 for storing and sharing the compressed swap pages. When swapping occurs by a lack of memory, a domainU calls memory callback functions in order to capture a swap page. Then the swap page is compressed by our compression function connected with the callback function. And the domain0 checks the page sharing table to find whether the same compressed swap page exist. If the compressed swap page already exists, the domain0 does not store new compressed swap page. Instead of the storing, the domain0 gives the address of the existing compressed swap page to the domainU. By this sharing step, the same compressed swap page is not stored into the domain0. It can reduce memory consumption of compressed swap pages in the domain0.

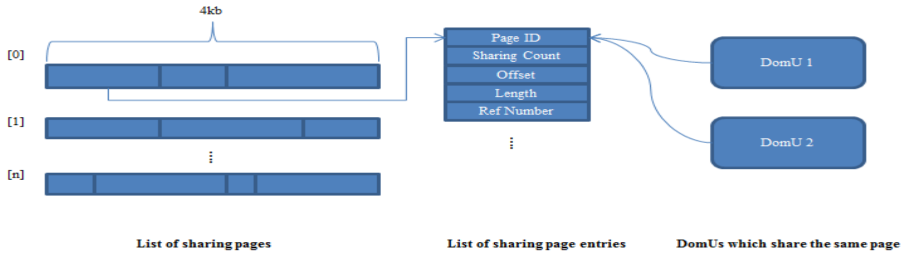

figure1. The page sharing table in the domain0

Figure1 shows the page sharing table in the domain0. That consists of two lists. The list of sharing pages has actual memories that contain compressed swap pages. And the list of sharing page entries maintains the information of compressed swap pages. Each entry contains page id, the number of sharing domains, length.

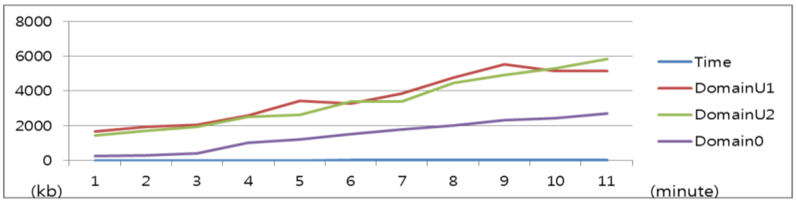

figure2. The amount of swap pages

Figure2 shows the amount of pages in each domain during compiling the kernel. The domainU1 made swap pages about 5000kb and the domainU2 made swap pages about 6000kb. However, the domain0's compressed swap pages are about 2700kb. As shown this result, our method reduces 75% of the amount of all swap pages by sharing swap pages. This means that our method can overcome the lack of memory caused by swap pages.

3. Conclusions

In the virtualization for the embedded systems, it is important to provide sufficient memory. Although Prior solution provides swapping for embedded systems using compressing swap pages and storing them into their memory, it cannot support enough memory for the applications. To solve this problem, this paper proposes compressing swap pages, storing compressed swap pages into another virtual machine with enough memory and sharing them between the virtual machines. As shown our experiment result, we reduce memory consumption of swap pages in the whole system.

Acknowledgement

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MEST) (No.2010-0029180) with KREONET.

References

[1] L. Yang et al., "Online Memory Compression for Embedded Systems," ACM Transactions on Embedded Computing Systems, Vol. 9, No. 3, February 2010.

[2] Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, Andrew Warfield, "Xen and the Art of Virtualization". SOSP, Oct. 2003.