

Memory Hotplug for Energy Savings of HPC systems*

Shinobu Miwa
miwa@is.uec.ac.jp

The University of Electro-Communications, Tokyo, Japan

Hiroki Honda
honda@is.uec.ac.jp


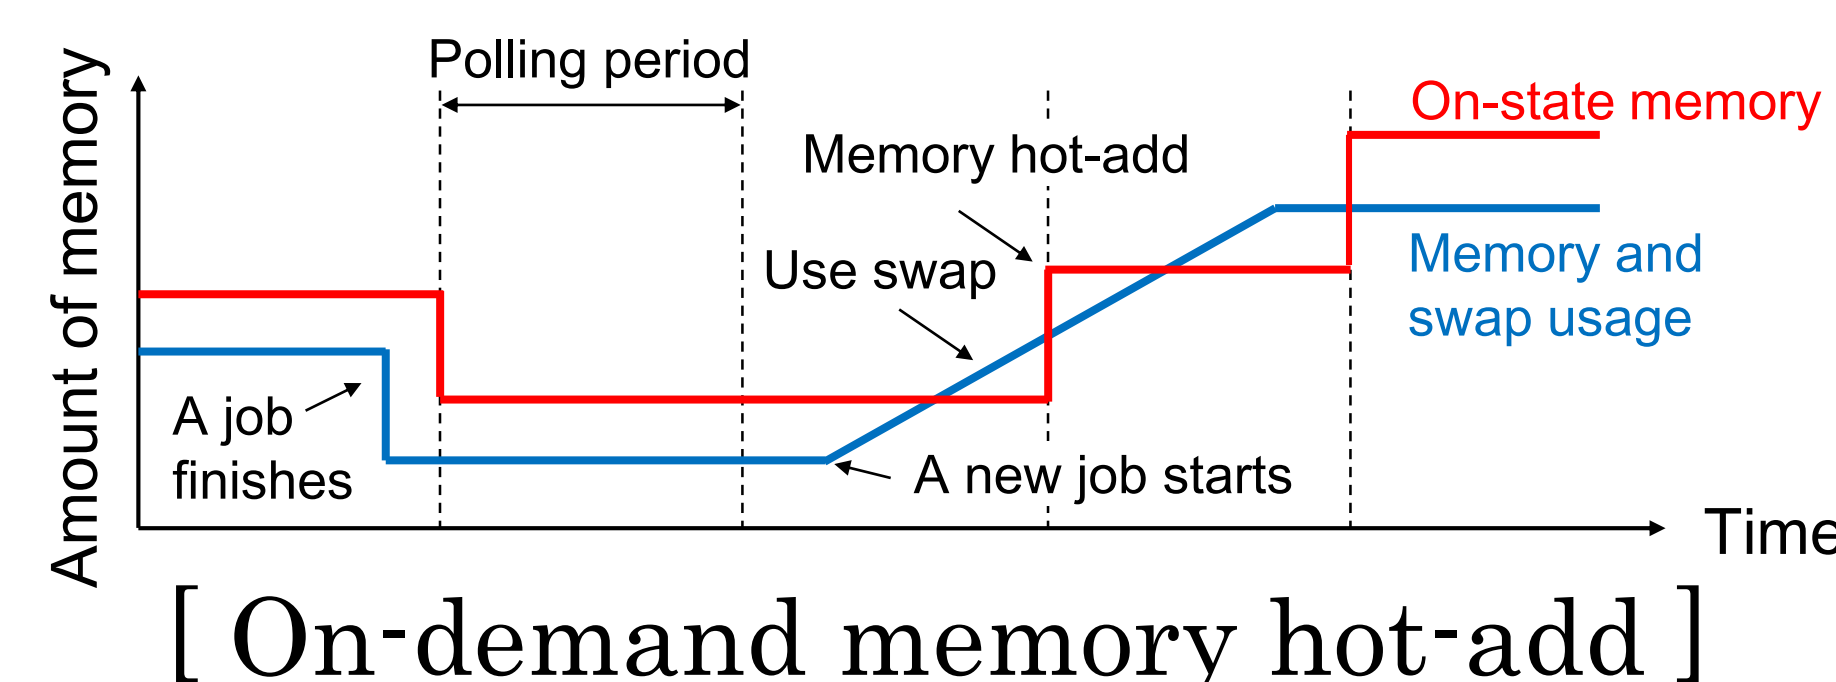

* This research is supported by JST under CREST

Background

- Energy savings of HPC systems is challenging for system developers
- Reduction of standby power is important for increased energy savings because many hardware devices in an HPC system are unused for long periods
- Many techniques have been proposed to reduce the standby power of CPUs, but there are a few techniques to reduce the standby power of memories

Energy Savings Using Memory Hotplug

- **Memory hotplug** is a technique to plug and pull DIMMs in a running system
 - Supported by Linux kernel 3.9 or later
 - Also require hardware support to physically plug and pull DIMMs
 - Usually used for increase of system availability
- Use memory hotplug for energy savings of HPC systems
 - Many memories are unused even during application running
 - An unused memory still consumes an amount of power (~ 1W per DIMM)
- On-demand memory hot-add
 - A daemon program running on a compute node turns off all unused memories when the node is idle
 - The daemon program polls memory and swap usage, and adds some memories when swapping happens
 - A polling period is a key to avoid thrashing while reducing the performance overhead of polling



Experiment

- Investigate the impact of polling periods of the daemon program on a server node running HPCC benchmarks
- Memory hotplug succeeded in saving memory size (**80% or more!**)
- The 0.1-second polling scheme shows little performance degradation (~ **2.7%**)

[Experimental system]

Name	Remarks
CPU	Xeon E5-2650 v2 (2.6GHz) x2 16 physical cores (32 logical cores)
Memory	16GB DDR3-1600 SDRAM x2 (32GB in total)
OS	Ubuntu 14.04 with Linux kernel 3.13

Summary and Future Work

- Proposed on-demand memory hot-add to save the memory standby power in HPC systems
- Memory hotplug has a great potential to save the memory standby power
- Will further investigate the impact of memory hotplug on the power and performance of HPC systems