

Working Set: Model and Characterization

Student: Ricardo Koller, PhD student, FIU

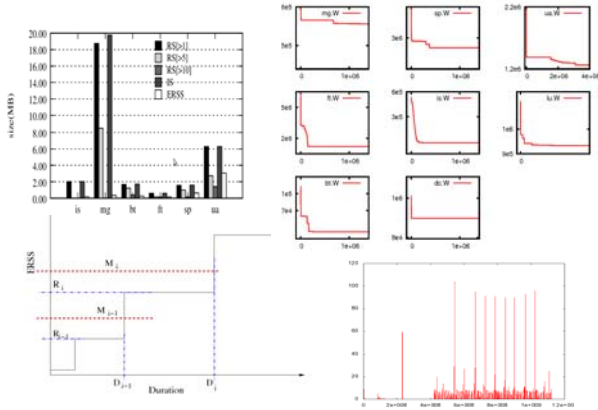
FIU Advisor: Raju Rangaswami, FIU

PIRE International Partner Advisor: Akshat Verma, IBM

I. Research Overview and Outcome

Model

- Inuse Set (IS):** An inuse set for a phase of duration θ_i is the subset of all distinct memory lines that are accessed in one phase.
- Reuse Set (RS):** The Reuse Set for a time window T is the subset of Inuse Set that is accessed more than once in the phase.
- Effective Reuse Set Size (ERSS):** The Effective Reuse Set Size is the minimum size of memory resources that need to be assigned to the program in that phase for it to not exhibit thrashing behavior.



ERSS Tree




Each node represents a phase. The parent-child relationship means that a phase includes a smaller phase.

References


- Peter J. Denning. Working sets past and present. In IEEE Tran. on Software Engineering, 1980.
- David K. Tam, Reza Azimi, Livio B. Soares, and Michael Stumm. Rapidmrc: Approximating l2 miss rate curves on commodity systems for online optimizations. In ASPLOS, 2009
- Ashutosh S. Dhodapkar and James E. Smith. Managing multi-configuration hardware via dynamic working set analysis. In Proc. ISCA, 2002

Characterization (tools implemented)

Probe based cache partitioning




Partitioning using a probe program that consumes a specific amount of cache. The only requirement for this probe is that it needs to access L2 lines faster than any other program. This can be achieved by using touch lines prefetching (TCB instruction on POWER6.)



Memory tracing on POWER6

Linux kernel patch that instruments the Performance Monitoring (PM) counter overflow interruptions and dumps the address of the load/store instructions from the SDAR. The PM is configured by the perf program.



L2 cache associativity based partitioning

The index portion of the address selects the set




The cache can be partitioned for an application if all the pages allocated to it are only touching some specific sets (for cache associativity).

Comparison of all methods


Technique	Accuracy	Intrusiveness	Overhead	Pre-requisites
Associativity	V. High	re-compilation kernel module	none	Access to aligned memory Limited Associativity
Probe	High	none	requires spare CPU	None
Trace	Medium	Access to hardware counters	1.3x	POWER5, POWER6 only
Simulator	Medium	none	10 to 100x	Standalone application only

Isolation

	ua	mg	dc	ua+mg	ua+dc
Misses(10 ⁶)	117	21	113	116	201
ERSS(MB)	3	0.7	0.5	3.85*	4.18*



II. International Experience



This was definitely a great trip. From one side, I had the pleasure to work with the same team and mentor, Akshat V., at IRL. And from the other side, I was able to visit one of the most interesting countries in the world for the second time. You can't get bored in India, and there is no way one can really get to know India in a couple of months.

III. Acknowledgement

The material presented in this poster is based upon the work supported by the National Science Foundation under Grant No. OISE-0730065. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.