The NUMA Problem

by Diego Esparza



What is NUMA?



Mapping pages? Copy-on-write and others...

Non-Uniform Memory Access (NUMA)





Invalidate copies when they are written to

So Many Policies...

Example Policies in Duke Study:

Wrm: Always migrate / replicate warm pages

LTDt: Defrost long frozen pages

... and 28 more policies

In many cases, page mapping policies hurt performance \rightarrow



Moving Threads Between Cores



Pages and threads can be mapped to cores simultaneously

Thread Communication Maps



← Static pattern means lots of thread interaction. Inevitably runs poorly on NUMA

Divide and Conquer

Dual Recursive Bipartitioning



Divide and Conquer

Dual Recursive Bipartition



Side Track: Asymmetric NUMA



Specifying splits



The machine - represented as a binary tree

Collecting Program Data



Divide and Conquer

Results



Resources

Larowe, R. and Schlatter Ellis, C., 1991. Experimental comparison of memory management policies for NUMA multiprocessors. ACM Transactions on Computer Systems, 9(4), pp.319-363.

Bolosky, Scott, Fitzgerald, Fowler and Cox, 1991. NUMA Policies and Their Relation to Memory Architecture. ASPLOS-IV Proceedings - Fourth International Conference on Architectural Support for Programming Languages and Operating Systems.

Lepers, Quéma, Fedorova, 2015. Thread and Memory Placement on NUMA Systems: Asymmetry Matters. Proceedings of the 2015 USENIX Annual Technical Conference.

Matthias Diener, Eduardo H.M. Cruz, Laércio L. Pilla, Fabrice Dupros, Philippe O.A. Navaux, 2015. Characterizing communication and page usage of parallel applications for thread and data mapping, Performance Evaluation, Volumes 88–89, Pages 18-36.

Pellegrini, 1994. Static Mapping by Dual Recursive Bipartitioning of Process and Architecture Graphs.