Data Partitioning on Heterogeneous Multicore Platforms

Ziming Zhong, Vladimir Rychkov, Alexey Lastovetsky
School of Computer Science and Informatics, University College Dublin
Belfield 4, Dublin, Ireland

ziming.zhong@ucdconnect.ie, {vladimir.rychkov, alexey.lastovetsky}@ucd.ie

Introduction

- Goal

  Load balancing MPI applications on heterogeneous multicore platforms.

- Techniques

  1) Inter-node data partitioning based on the nodes' Multicore Functional Performance Model.
  2) Intra-node data partitioning based on dynamic evaluation of performance of processing cores.

Multicore Functional Performance Model

Experiment testbed:

Host	Paradent	Paramount	Parapide	Parapluie
Processor	Xeon L5420	Xeon 5418 LV	Xeon X5570	Opteron 6164HE
CPU Speed	2.50 GHz	2.33 GHz	2.93 GHz	1.70 GHz
Num of CPUs	2	2	2	2
Num of Cores	8	4	8	24
L1 Cache	32KB	32KB	32KB	64KB
L2 Cache	6144KB	4096KB	256KB	512KB
L3 Cache	-	-	8MB	10MB
RAM	31GB	8GB	24GB	48GB
NUMA	NO	NO	YES	YES
Memory Band-width	1.33GT/s	1.33GT/s	6.4GT/s	6.4GT/s

Multicore Functional Performance Model -Definition:

$s(x) = \frac{x}{t(x)}$ is a positive continuous function representing speed of the node.

$x = \sum_{i=0}^{c-1} x_i$ is the total problem size processed on the node x_i is problem size assigned to individual core.

$t(x) = \max_{i=0}^{c-1} t_i(x_i)$ is the execution time of the parallel routine on the node, $t_i(x_i)$ is execution time on i-th core.

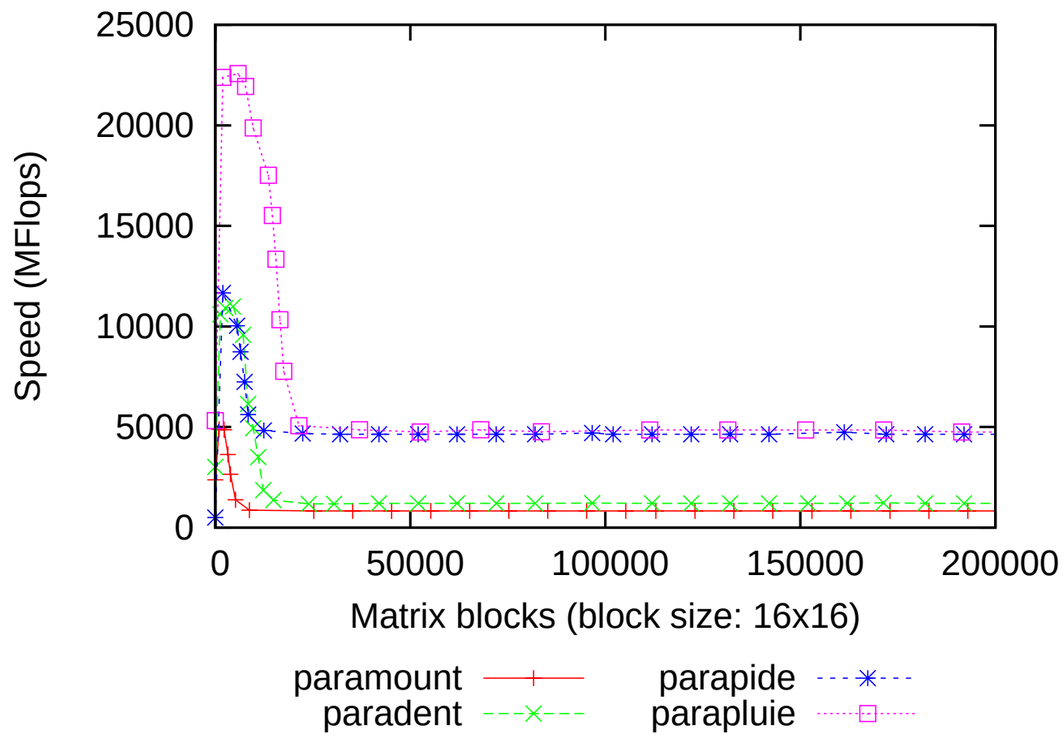

Inter-node Data Partitioning

Optimal data distribution based on Functional Performance Model [1]

$$\frac{x_1}{s_1(x_1)} = \frac{x_2}{s_2(x_2)} = \ldots = \frac{x_p}{s_p(x_p)}$$

$$x_1 + x_2 + \ldots + x_p = n$$

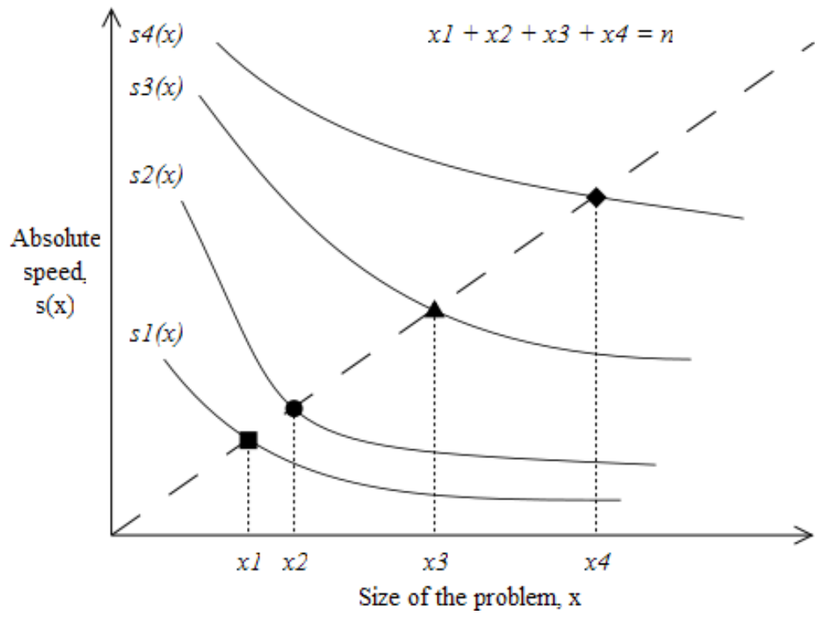

Results and Comparison

1. Results based on Multicore Functional Performance Model (MFPM)
2. Results based on Constant Performance Model (CPM)
3. Homogeneous data partitioning

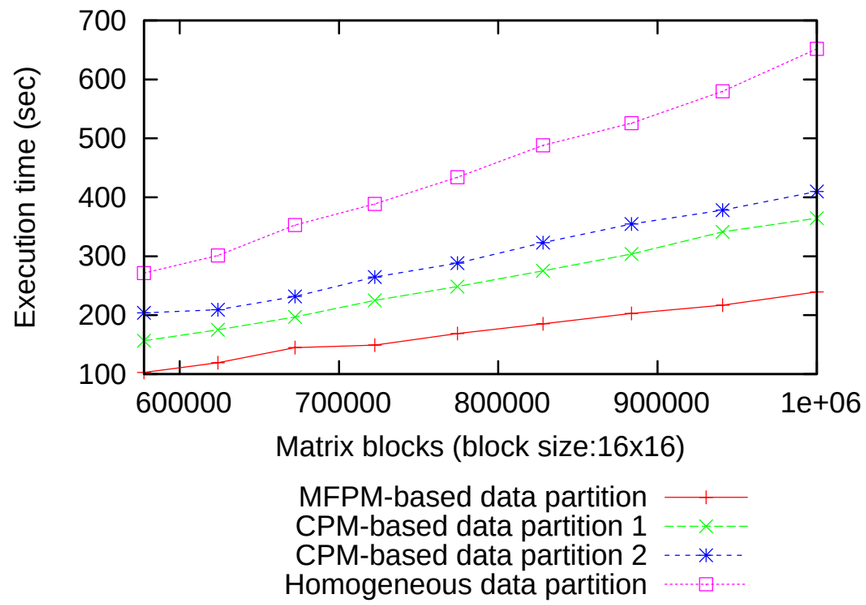

Intra-node Data Partitioning

Platform

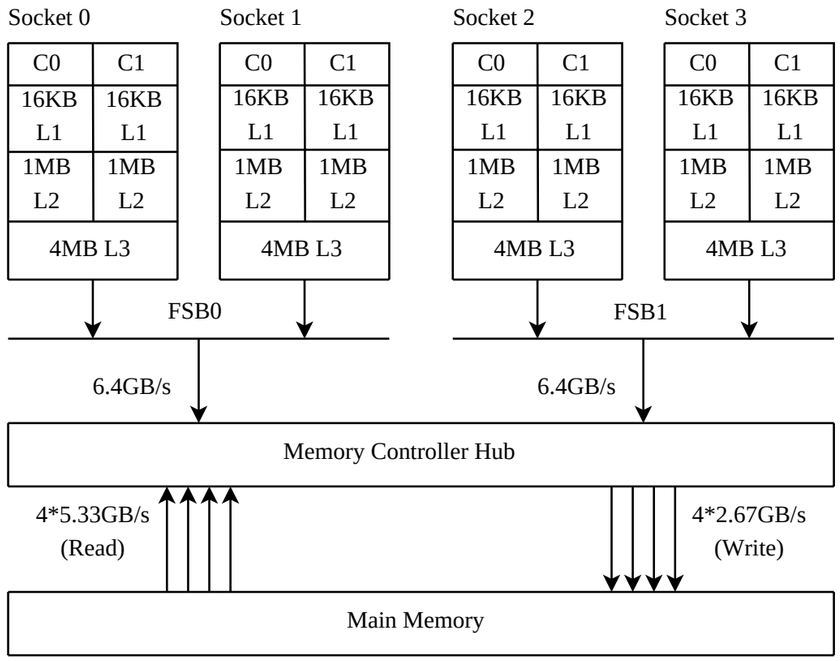

Heterogeneous Setup

- Socket 0: one process per core
- Socket 1-3: one process per socket

Speed

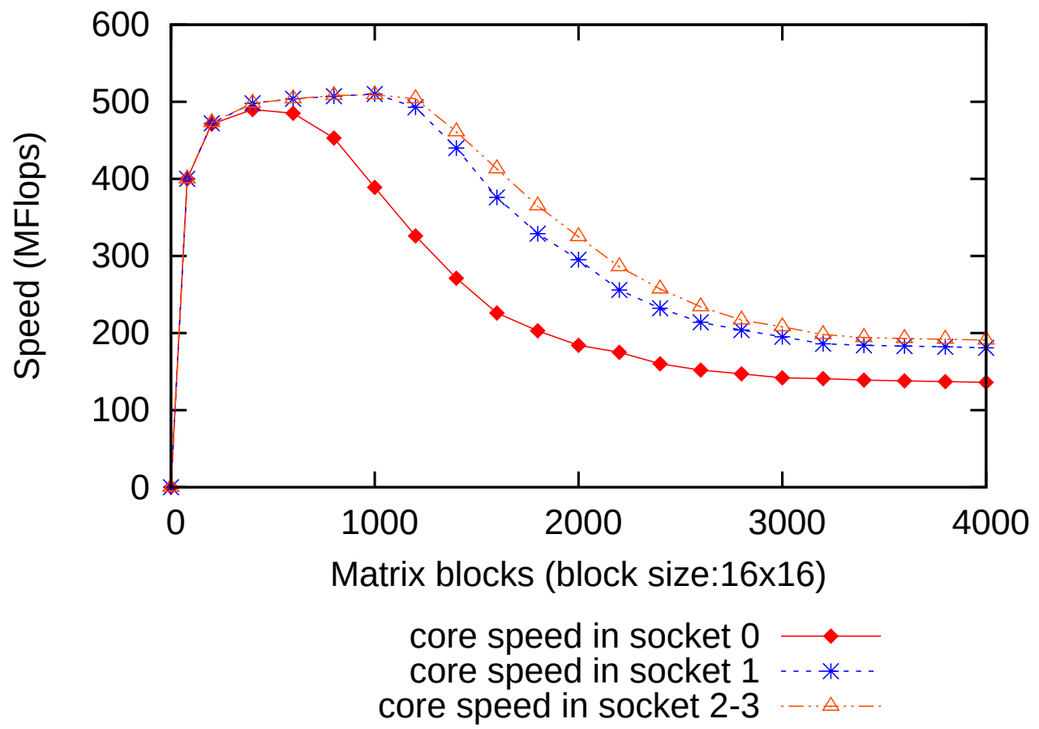

Two-step data partitioning

- Data are evenly distributed on cores to measure speed of each core.
- Data are redistributed in proportion to the speed of each core.

- Results of Intra-node data partitioning

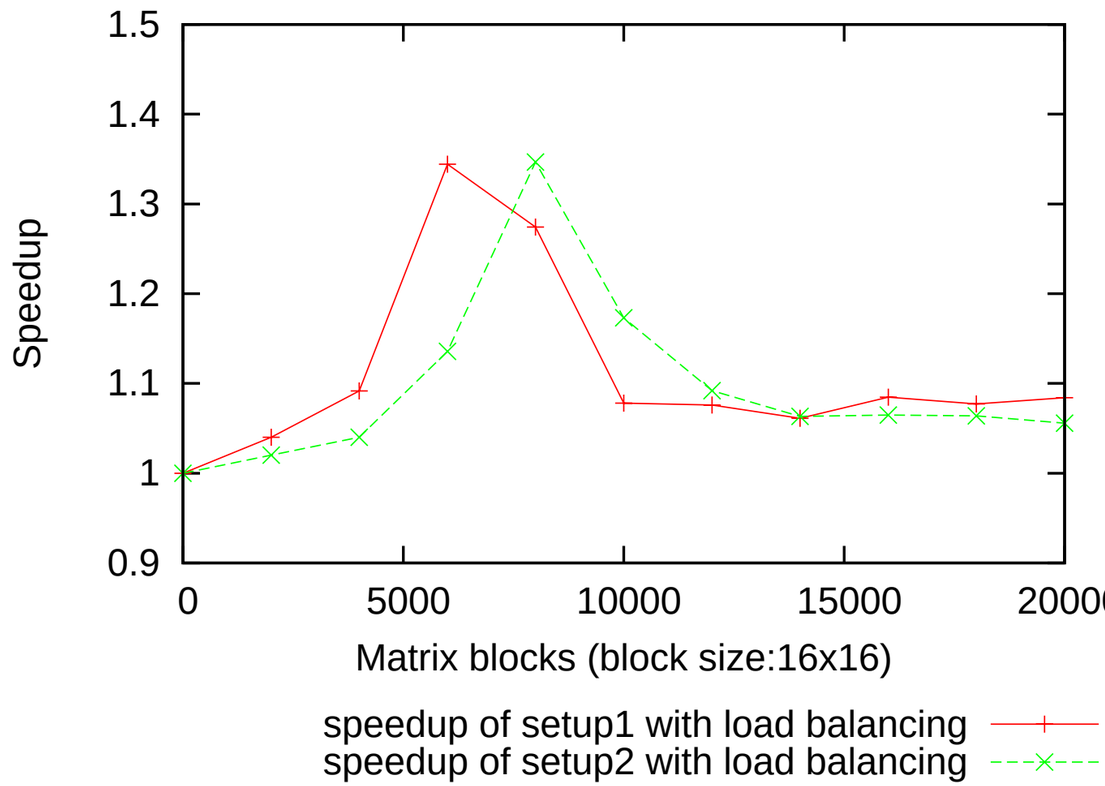

- Reference

[1] A. Lastovetsky, and R. Reddy, “Data Partitioning with a Functional Performance Model of Heterogeneous Processors,” *Int. J. High Perform. Comput. Appl.*, vol. 21, pp. 76-90, Feb. 2007.