

Thread Tailor

Dynamically Weaving Threads Together for
Efficient, Adaptive Parallel Applications

Janghaeng Lee, Haicheng Wu,
Madhumitha Ravichandran, Nathan Clark

**Georgia
Tech**



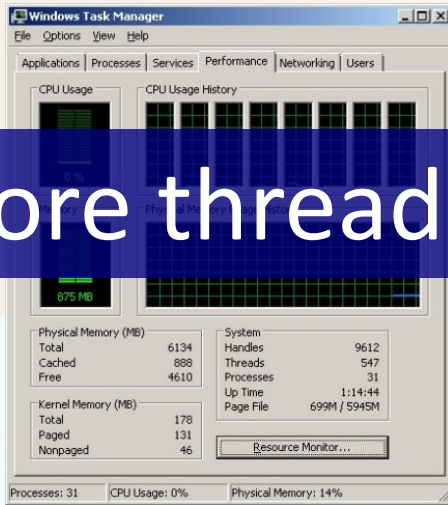
College of
Computing



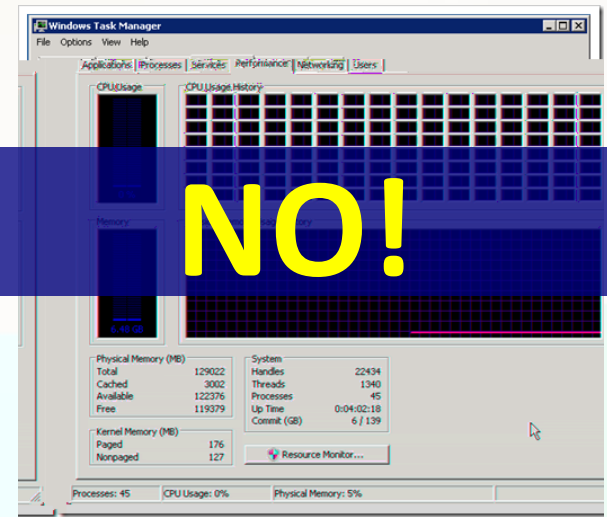
Motivation

- Hardware Trends
 - Put more cores in a single chip

More threads always win? **NO!**



2009



201X

- CPU intensive programs
 - Exploits Thread Level Parallelism



Optimal Number of Threads

- Too many threads
 - More synchronization
 - More contention for system resources
- Too few threads
 - Resource underutilization
- Who can decide the number?
 - **Not a programmer**



Why NOT?

- Input changes
 - Various working-set size

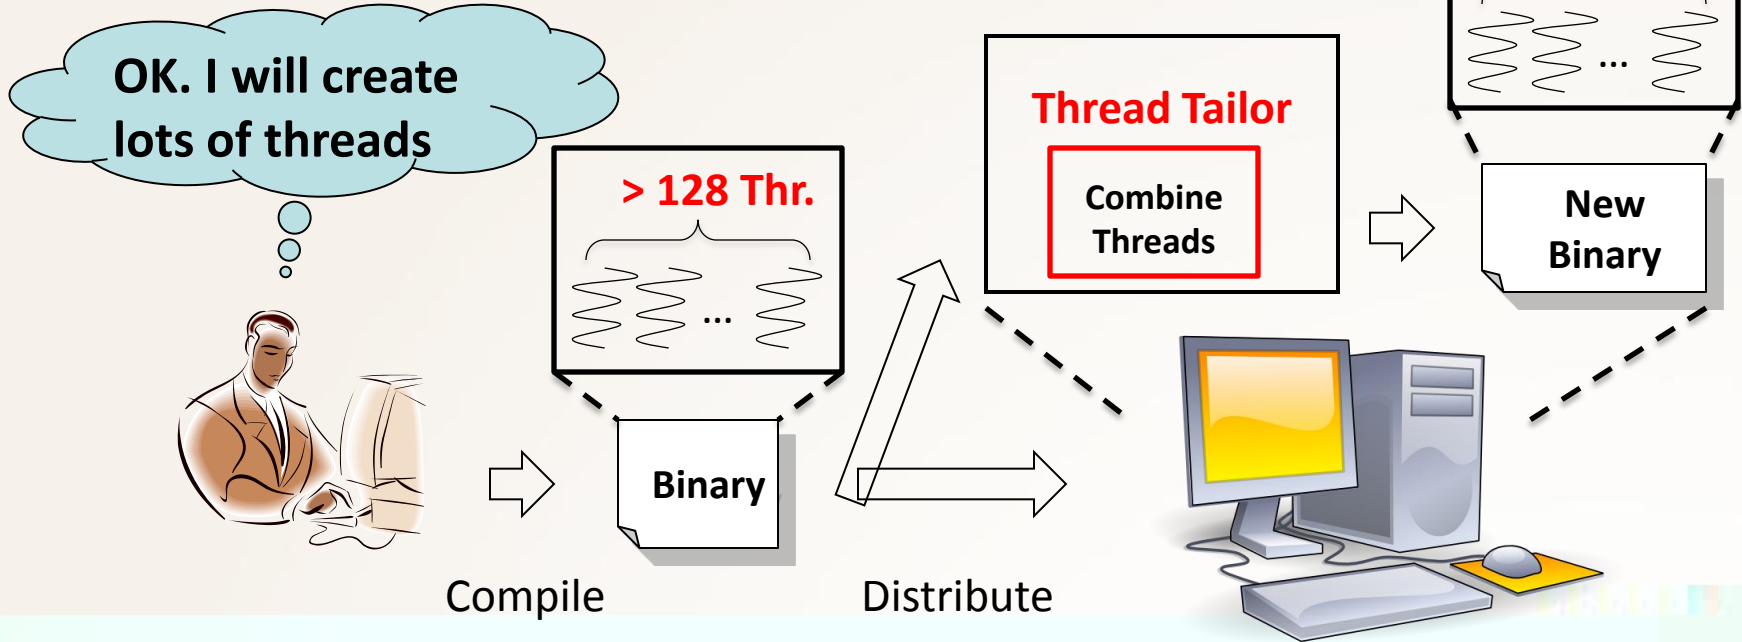
- The system changes

Decision must be made at runtime

- Various available resources

- Hardware changes
 - Various L2/L3 cache structure / size, etc.

Proposal

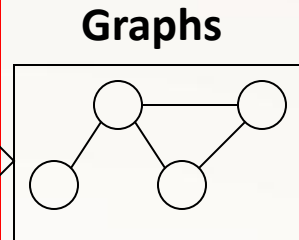
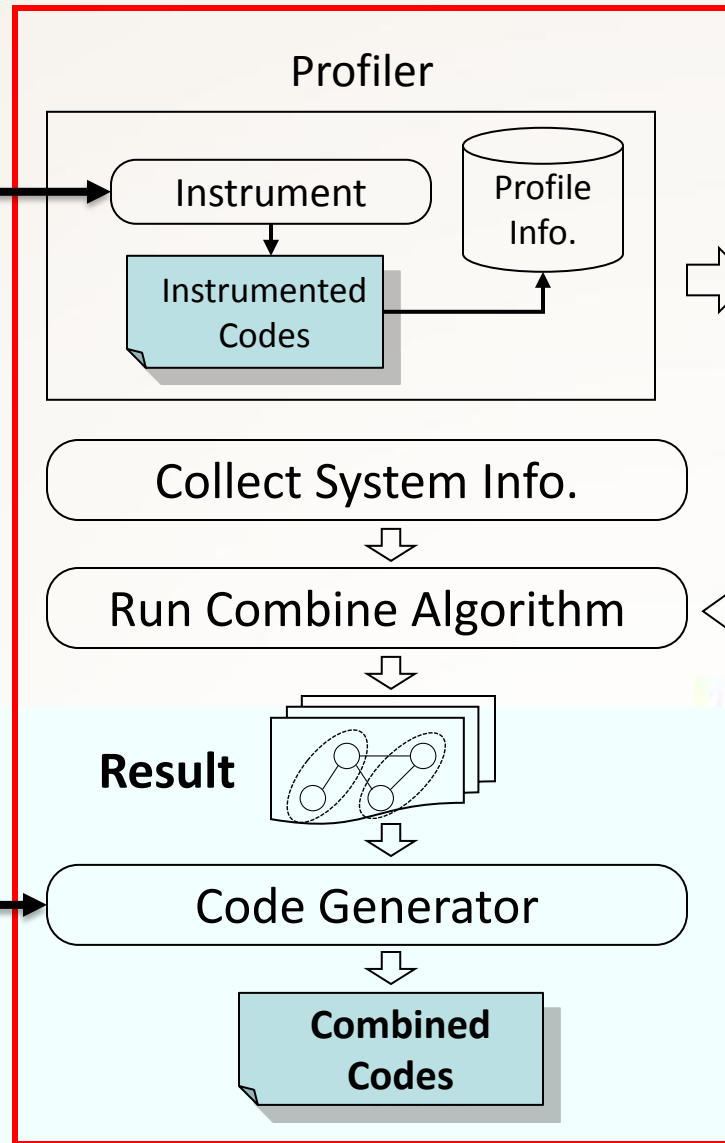
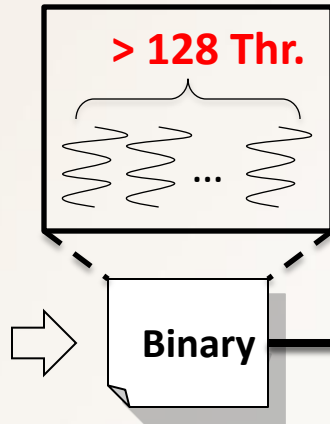


- Combining Threads

- Group Several Threads into a Single Thread

- Threads in the same group are executed in serial
- Executed on the SAME core

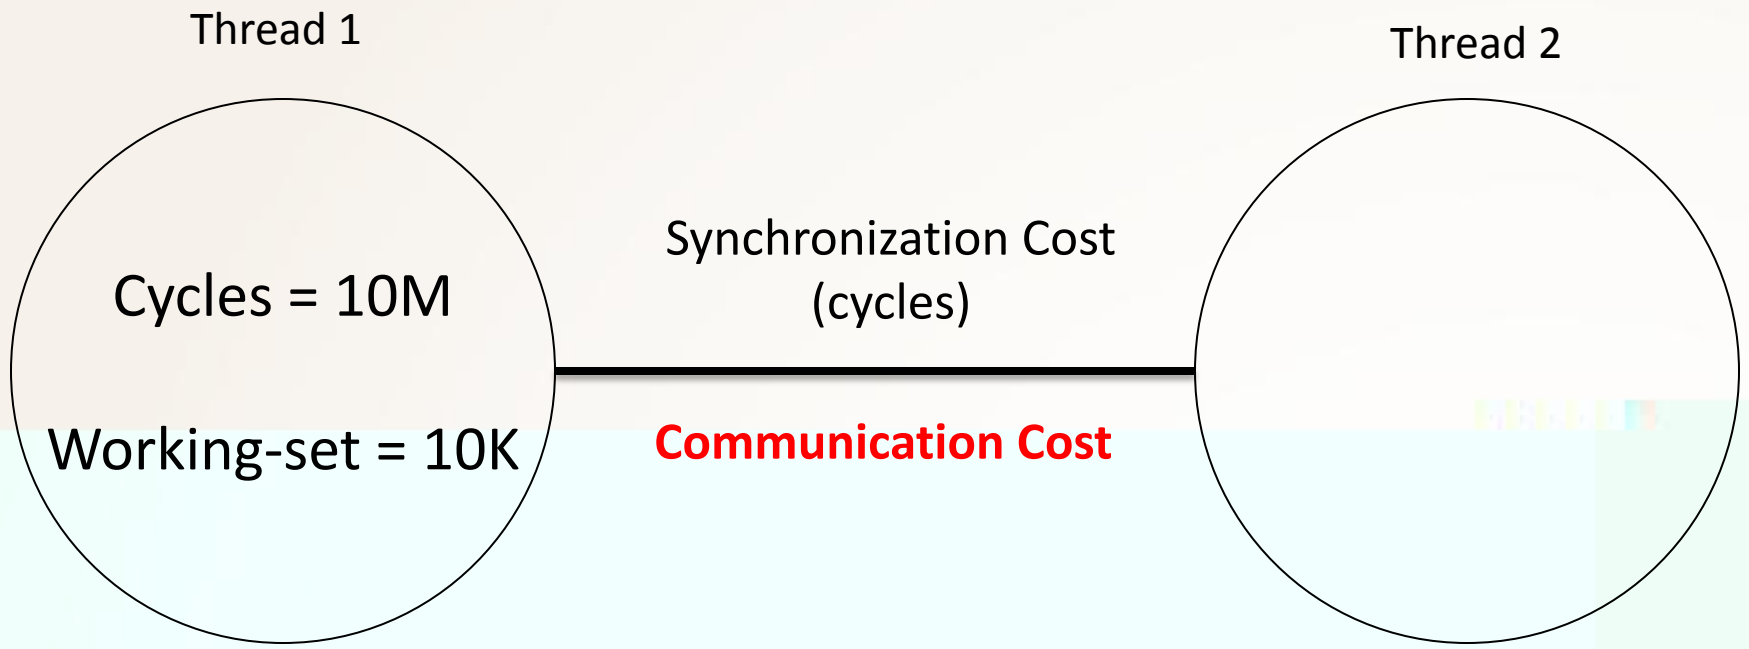
Details



Development | Distribution



Graph Construction





Communication Cost

- Intuition : STORE Instruction causes coherence miss in cache
- Log Memory Access per Thread

Thread 1		
Address	LD Count	ST Count
...
0x00001234	5	10
0x00001338	4	9
...
0x00004000	7	7
...

Thread 2		
Address	LD Count	ST Count
...
0x00001234	0	7
0x00002000	4	4
...
0x00004000	3	8
...

LD LD

ST ST

Graph

① — 29 — ②

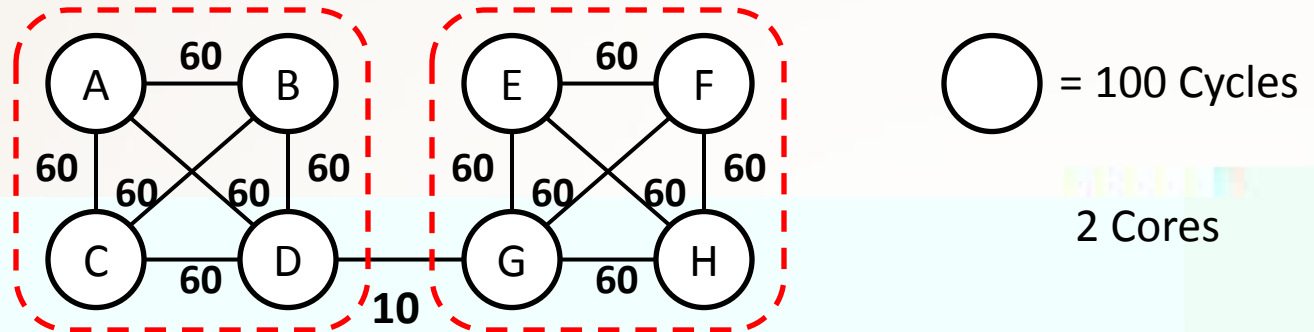
$$0x00001234: \text{MIN}(5, 7) + \text{MIN}(10, 0) + \text{MIN}(10, 7) = 12$$

$$0x00004000: \text{MIN}(7, 8) + \text{MIN}(7, 3) + \text{MIN}(7, 8) = 17$$

$$\text{Total Communication Cost: } 12 + 17 = \mathbf{29}$$

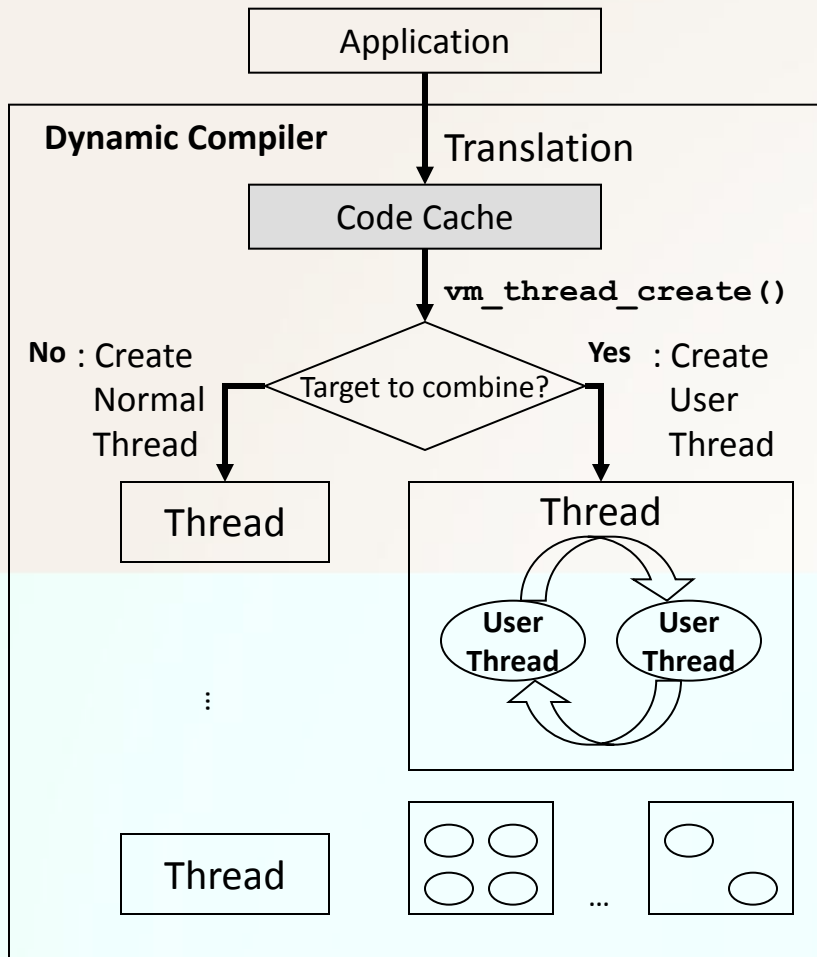
Combining Algorithm

- Kernighan-Lin(KL) Graph Partitioning Heuristic
 - Goal : Minimize Execution Cycles
 - Precondition : Combined Threads \leq Cores



Partition 1	Partition 2	Partition 1 Cycle Estimation	Partition 2 Cycle Estimation	Move From	Move Node
BCDG	AEFH	210	220	2	A
ABCDG	EEH	130	120	1	G
ABCD	EFGH	40	40	1	D

Thread Combining



Replace Thread APIs with Wrapper Functions

Wrapper Function for Thread Creation

Context Switched by Dynamic Compiler

Serially Execute User Threads in Real Thread



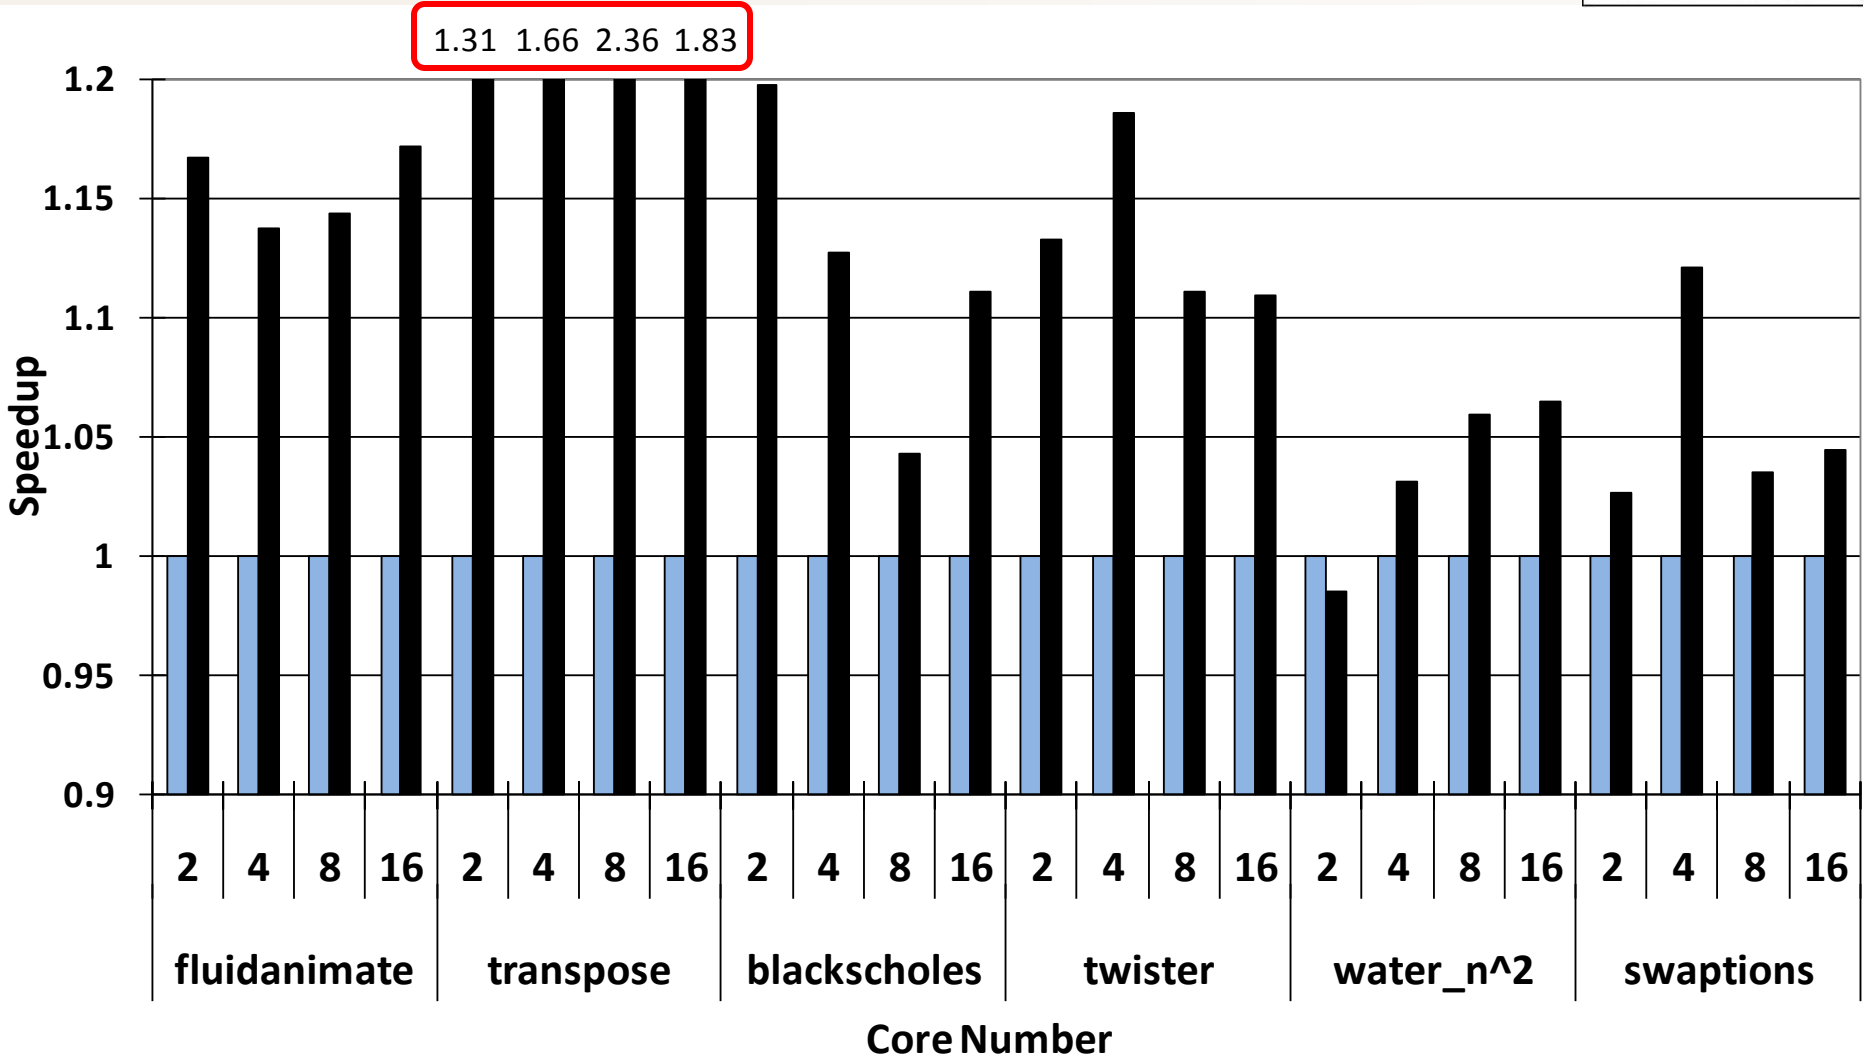
Experimental Setup

- 2 cores
 - Intel Core 2 Duo 6600 (2.4 Ghz)
- 4 cores
 - Intel Core 2 Quad Q6600 (2.4.Ghz)
- 8 cores
 - 2 Quad-core CPUs with SMT
 - Intel Xeon E5520 (2.26 Ghz)
- 16 cores (Logical)
 - 2 Quad-core CPUs with SMT and HyperThreading
 - Intel Xeon E5520 (2.26 Ghz)

Results



■ #thr=#core
■ Thread Tailor





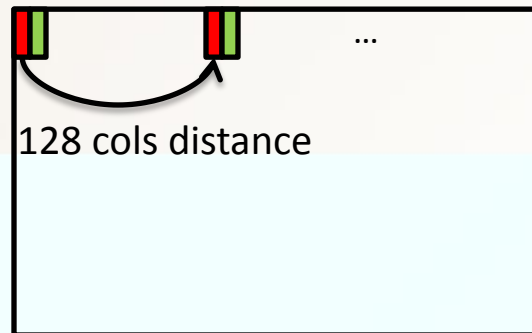
Result Analysis - Transpose

- Transpose $m * n$ matrix to $n * m$

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \Rightarrow \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$

- Parallel Transpose

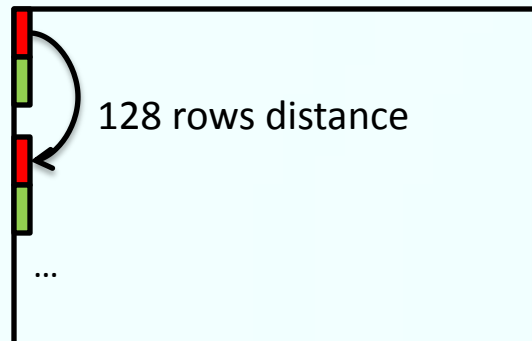
Input Matrix



 Thread 1

 Thread 2

Output Matrix





Result Analysis - Transpose

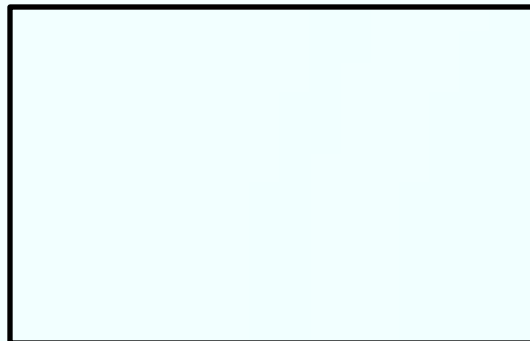
- Transpose $m * n$ matrix to $n * m$

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \Rightarrow \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$

Input Matrix
16K x 16K



Output Matrix
16K x 16K



Intel Nehalem

Core 0

L1 private (32K)

L2 private (256K)

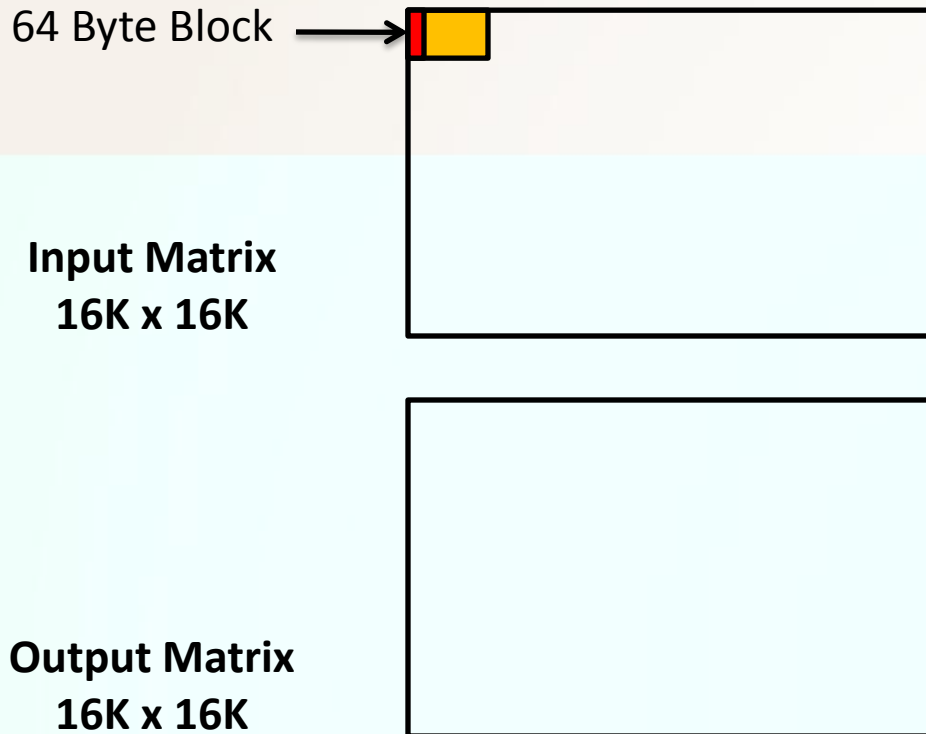
L3 Shared (8M)



Result Analysis - Transpose

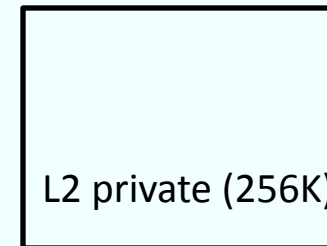
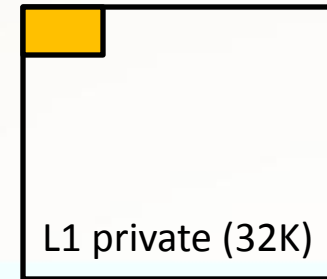
- Transpose $m * n$ matrix to $n * m$

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \Rightarrow \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$



Intel Nehalem

Core 0





Result Analysis - Transpose

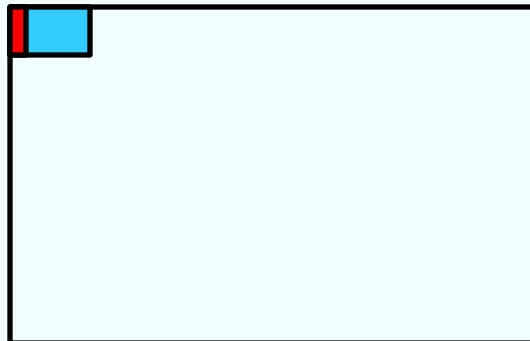
- Transpose $m * n$ matrix to $n * m$

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \Rightarrow \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$

Input Matrix
16K x 16K

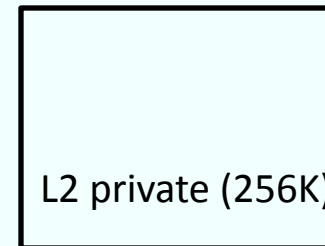
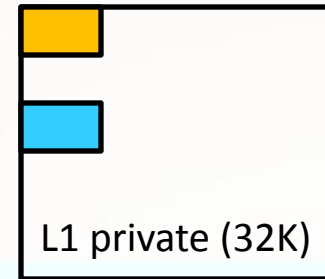


Output Matrix
16K x 16K



Intel Nehalem

Core 0



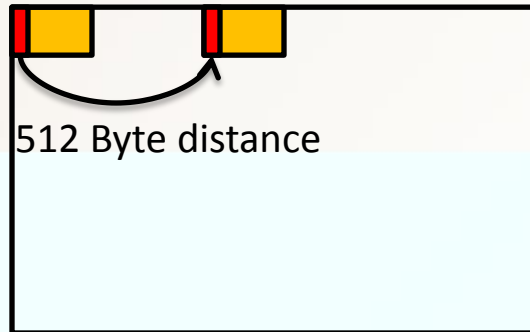


Result Analysis - Transpose

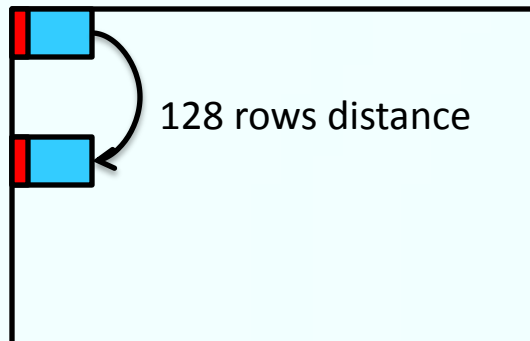
- Transpose $m * n$ matrix to $n * m$

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \Rightarrow \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$

Input Matrix
16K x 16K

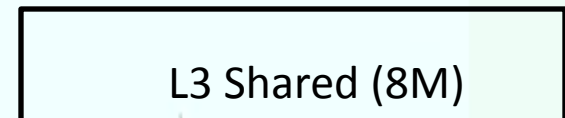
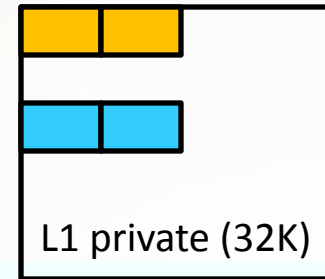


Output Matrix
16K x 16K



Intel Nehalem

Core 0

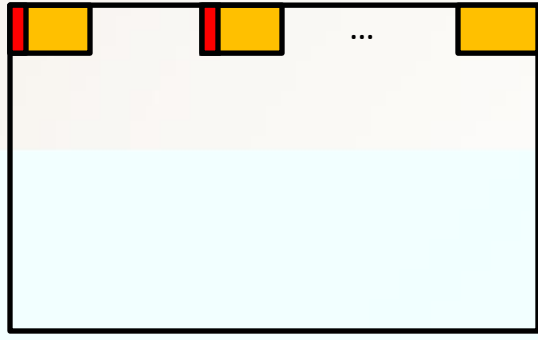




Result Analysis - Transpose

- Transpose $m * n$ matrix to $n * m$

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \Rightarrow \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$



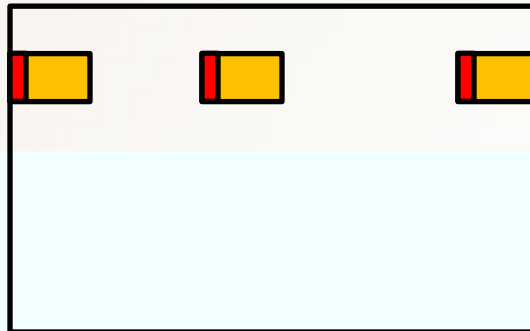


Result Analysis - Transpose

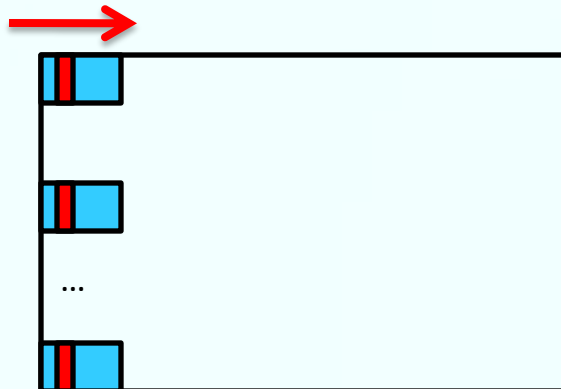
- Transpose $m * n$ matrix to $n * m$

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \Rightarrow \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$

Input Matrix
16K x 16K

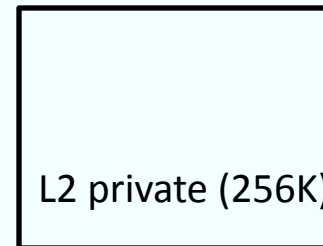
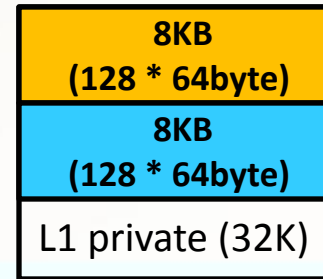


Output Matrix
16K x 16K



Intel Nehalem

Core 0



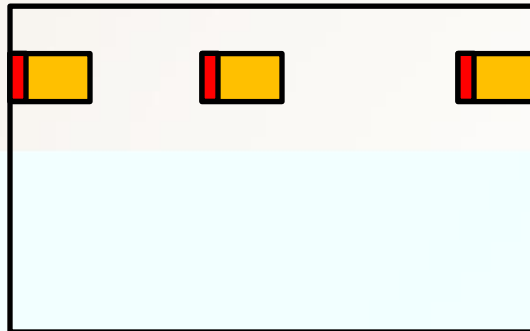


Result Analysis - Transpose

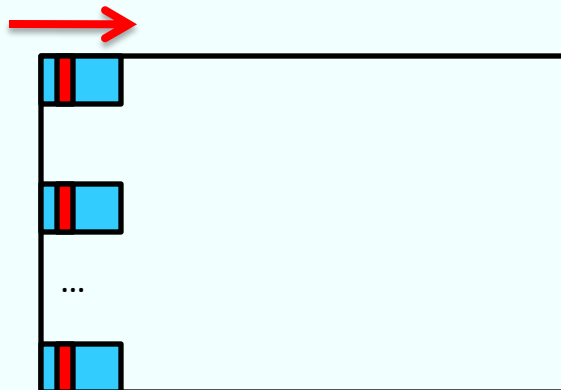
- Transpose $m * n$ matrix to $n * m$

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \Rightarrow \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$

Input Matrix
16K x 16K



Output Matrix
16K x 16K



Intel Nehalem

Core 0

8KB
(128 * 64byte)

8KB
(128 * 64byte)

WRITE HIT!

L1 private (32K)

...

L2 private (256K)

L3 Shared (8M)



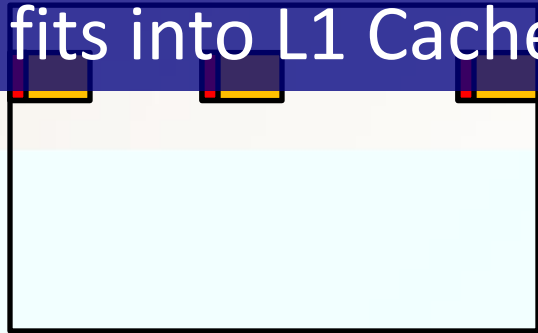
Result Analysis - Transpose

- Transpose $m * n$ matrix to $n * m$

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \Rightarrow \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$

Working-set fits into L1 Cache (No Capacity Miss!)

Input Matrix
16K x 16K

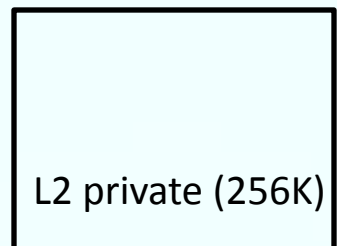
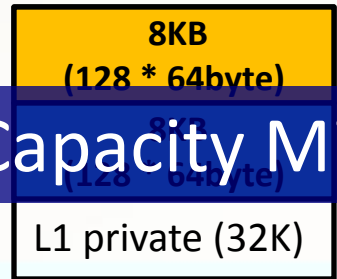


Output Matrix
16K x 16K



Intel Nehalem

Core 0





Summary

- Choosing Optimal Number of Threads is Hard
- Thread Tailor Ease the Pain
 - Graph Representation
 - Combine Threads at Runtime



Thank you