

Lazy Parallel Kronecker Algebra-operations on Heterogeneous Multicores

EuroPar 2017

Wasuwee Sodsong¹, Robert Mittermayr², Yoojin Park¹,
Bernd Burgstaller¹ and Johann Blieberger²

¹Yonsei University, Seoul, Korea

²Vienna University of Technologies, Vienna, Austria






Embedded Systems Languages and Compilers Lab, Yonsei University
Automation System Group, Vienna University of Technologies

Motivation: Computing Thread Interleavings

2

- *Arbitrary* interleavings of threads such that...
 - ▣ the order on computation steps is taken from threads' program order



Interleavings

a · b · c · d

a · c · b · d

a · c · d · b

c · d · a · b

c · a · d · b

c · a · b · d


- Totality of all possible thread interleavings is well-suited for concurrent program analysis

Synchronization Primitives

3

- Semantics of synchronization primitives constrain possible interleavings
 - ▣ Not all interleavings are **valid**

Example: binary semaphore



Interleavings

$p \cdot v \cdot p \cdot v$ ✓

$p \cdot p \cdot v \cdot v$ ✗

$p \cdot p \cdot v \cdot v$ ✗

$p \cdot v \cdot p \cdot v$ ✓

$p \cdot p \cdot v \cdot v$ ✗


$p \cdot p \cdot v \cdot v$ ✗

Synchronization Primitives

4

- Semantics of synchronization primitives allow additional interleavings that constitute deadlocks

Example: **self-deadlock** on binary semaphore




Interleavings

p

$p \cdot v \cdot p$


- Behavior of semaphores can be modelled by a deterministic finite automaton (DFA)



Encoding of Threads as Adjacency Matrices

5

- DFA representation of a thread or a semaphore can be encoded as an adjacency matrix
 - ▣ Rows represent node IDs
 - ▣ Columns represent successor IDs





Kronecker Product: Lock-step Execution of Threads


6

- **Kronecker product:** Given an m-by-n matrix A and a p-by-q matrix B,

$$A \otimes B = \begin{pmatrix} a_{1,1} \cdot B & \cdots & a_{1,n} \cdot B \\ \vdots & \ddots & \vdots \\ a_{m,1} \cdot B & \cdots & a_{m,n} \cdot B \end{pmatrix}$$



$$A \otimes B = \begin{pmatrix} 0 & 0 & bc & bd \\ 0 & 0 & 0 & 0 \\ 0 & 0 & ac & ad \\ 0 & 0 & 0 & 0 \end{pmatrix}$$





Kronecker Product: Lock-step Execution of Threads


7

- **Kronecker product:** Given an m-by-n matrix A and a p-by-q matrix B,

$$A \otimes B = \begin{pmatrix} a_{1,1} \cdot B & \cdots & a_{1,n} \cdot B \\ \vdots & \ddots & \vdots \\ a_{m,1} \cdot B & \cdots & a_{m,n} \cdot B \end{pmatrix}$$




- Executed threads A and B in lock-step
- \otimes can be used to **synchronize** threads



Synchronization: Lock-step Execution of Threads and Semaphores


8

Thread



$$t = \begin{pmatrix} 0 & p & 0 \\ 0 & 0 & v \\ 0 & 0 & 0 \end{pmatrix}$$

Semaphore p



$$s = \begin{pmatrix} 0 & p \\ v & 0 \end{pmatrix}$$

$$t \otimes s = \begin{pmatrix} 0 & 0 & 0 & pp & 0 & 0 \\ 0 & 0 & pv & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & vp \\ 0 & 0 & 0 & 0 & vv & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & p & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & v & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$




□ For simplicity we write $x \cdot x = x$ and $x \cdot y = 0$

Kronecker Sum: Interleaved Execution of Threads


9

- **Kronecker sum:** Given a square matrix A of order m and B of order n ,

$$A \oplus B = A \otimes I_n + I_m \otimes B.$$




$$A = \begin{pmatrix} 0 & a & 0 \\ 0 & 0 & b \\ 0 & 0 & 0 \end{pmatrix}$$



$$B = \begin{pmatrix} 0 & c & 0 \\ 0 & 0 & d \\ 0 & 0 & 0 \end{pmatrix}$$

$$A \oplus B = \begin{pmatrix} 0 & c & 0 & a & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & d & 0 & a & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & a & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & c & 0 & b & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & d & 0 & b & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & b \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & c & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & d \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$




Kronecker Sum: Interleaved Execution of Threads


10

- **Kronecker sum:** Given a square matrix A of order m and B of order n ,

$$A \oplus B = A \otimes I_n + I_m \otimes B.$$




$$A = \begin{pmatrix} 0 & a & 0 \\ 0 & 0 & b \\ 0 & 0 & 0 \end{pmatrix}$$



$$B = \begin{pmatrix} 0 & c & 0 \\ 0 & 0 & d \\ 0 & 0 & 0 \end{pmatrix}$$

- All thread **interleavings**
- \oplus can be used to model concurrency




Concurrent Threads and Semaphores


11

- Encode threads' control flow graphs and synchronization primitives as adjacency matrices.

Control flow graph of thread $t^{(i)}$



Semaphore $s^{(j)}$




- $T = \bigoplus_{i=1}^k t^{(i)}$ models all interleavings of the threads
- $S = \bigoplus_{j=1}^r s^{(j)}$ models all interleavings of the semaphores

Synchronizing Threads and Semaphores

12


- Behaviors of overall systems can be modelled by




- T_s contains only semaphore calls
- T_v contains the other edge labels
- $T = T_s + T_v$

Example: Overall System

13



$$\text{Thread } t = \begin{pmatrix} 0 & p & p & 0 \\ 0 & 0 & 0 & \textcolor{red}{p} \\ 0 & 0 & 0 & v \\ 0 & 0 & 0 & 0 \end{pmatrix}$$




$$\text{Semaphore } s = \begin{pmatrix} 0 & p \\ v & 0 \end{pmatrix}$$

$$P = T_s \otimes S + T_v \oplus S$$


$$P = t \otimes s = \begin{pmatrix} 0 & p & p & 0 \\ 0 & 0 & 0 & \textcolor{red}{p} \\ 0 & 0 & 0 & v \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes \begin{pmatrix} 0 & p \\ v & 0 \end{pmatrix} =$$

Example: Overall System

14



$$\text{Thread } t = \begin{pmatrix} 0 & p & p & 0 \\ 0 & 0 & 0 & \textcolor{red}{p} \\ 0 & 0 & 0 & v \\ 0 & 0 & 0 & 0 \end{pmatrix}$$



$$P = T_s \otimes S + T_v \oplus S$$

$$P = t \otimes s = \begin{pmatrix} 0 & p & p & 0 \\ 0 & 0 & 0 & \textcolor{red}{p} \\ 0 & 0 & 0 & v \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes \begin{pmatrix} 0 & p \\ v & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & p & 0 & p & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \textcolor{red}{p} \\ \textcolor{red}{0} & \textcolor{red}{0} & \textcolor{red}{0} & \textcolor{red}{0} & \textcolor{red}{0} & \textcolor{red}{0} & \textcolor{red}{0} & \textcolor{red}{0} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & v & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \end{matrix}$$

Problems

15

- Combination of **Kronecker algebra** operations generates **all possible** (arbitrary) thread interleavings subject to semantics of synchronization primitives
- The number of interleavings increases **exponentially** in the number of threads
 - The state explosion problem
 - **But:** Not all interleavings need to be computed!

Contributions

16

1. Two-step lazy evaluation scheme
 - ▣ Construct expression tree
 - ▣ Lazily evaluate *relevant* thread interleavings
2. Kronecker algebra operations optimized for multicore CPUs
3. Execution scheme that utilizes both the multicore CPU and the GPU
 - ▣ GPU: conducts lazy evaluation.
 - ▣ CPU: maintains the computed thread interleavings and coordinate the GPU-based evaluation process
4. Introduce a fast operation on a special case of Kronecker sums


Lazy Evaluation

17

Observations

- Adjacency matrices are sparse
- Not all nodes are reachable from the start node
- It is not necessary to compute the entire matrix


$$t \otimes s = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \end{matrix} \\ \begin{pmatrix} 0 & 0 & 0 & \mathbf{p} & 0 & \mathbf{p} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{p} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \mathbf{v} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} & \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \end{matrix} \end{matrix}$$



Lazy Evaluation

18

- We never compute an entire matrix
 - ▣ Instead, we compute only the **non-zero** elements of the reachable nodes



- **Step 1:** Construct an expression tree
 - ▣ Leaf nodes (operands) are stored as *sparse matrices*
 - ▣ Internal nodes (operators) are stored as *lazy matrices*
 - the algebra operator and pointers to the operands
 - ▣ No operator evaluation yet!

Lazy Evaluation



19

- **Step 2:** Lazily evaluate Kronecker algebra operations
 1. Find successors of the start node by evaluating the expression tree
 2. If a successor has not been processed before, insert it to a work-queue
 3. Repeatedly process nodes from the work-queue until it is empty


Example: self-deadlock thread

$$t \otimes s = \begin{pmatrix} 0 & p & p & 0 \\ 0 & 0 & 0 & p \\ 0 & 0 & 0 & v \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes \begin{pmatrix} 0 & p \\ v & 0 \end{pmatrix} =$$


$$\begin{pmatrix} 0 & 0 & 0 & p & 0 & p & 0 & 0 \\ ? & ? & ? & ? & ? & ? & ? & ? \\ ? & ? & ? & ? & ? & ? & ? & ? \end{pmatrix}$$



work-queue



Processed nodes



Semaphore's Kronecker Sum Optimization

20

- We can calculate a **series** of Kronecker sums of same-sized matrices in one step.
 - ▣ a node ID
 - ▣ number of Kronecker sum operations

$$s_1 =$$

0	p ₁
v ₁	0

$$s_1 \oplus s_2 =$$

0	p ₂	p ₁	0
v ₂	0	0	p ₁
v ₁	0	0	p ₂
0	v ₁	v ₂	0


$$s_1 \oplus s_2 \oplus \mathbf{s}_3 =$$

0	p₃	p ₂	0	p ₁	0	0	0
v₃	0	0	p ₂	0	p ₁	0	0
v ₂	0	0	p₃	0	0	p ₁	0
0	v ₂	v₃	0	0	0	0	p ₁
v ₁	0	0	0	0	p₃	p ₂	0
0	v ₁	0	0	v₃	0	0	p ₂
0	0	v ₁	0	v ₂	0	0	p₃
0	0	0	v ₁	0	v ₂	v₃	0

Kronecker Algebra on a Multicore CPU

21

- ❑ A thread finds all successors of a given node
- ❑ Employ a **hash-function** which hashes the node IDs of the successors
 - ❑ The hash-function guarantees **one-to-one** assignment of node IDs to worker threads
- ❑ Each worker thread maintains a local hash-table of processed nodes



$$t \otimes s = \begin{pmatrix} 0 & p & p & 0 \\ 0 & 0 & 0 & p \\ 0 & 0 & 0 & v \\ 0 & 0 & 0 & 0 \end{pmatrix} \otimes \begin{pmatrix} 0 & p \\ v & 0 \end{pmatrix} =$$

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ 0 & 0 & 0 & p & 0 & p & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & v & 0 \\ ? & ? & ? & ? & ? & ? & ? & ? \end{pmatrix}$$

Kronecker Algebra using a CPU+GPU

22

- A GPU has relatively small memory compared to a CPU
 - ▣ unable to keep the continuously growing list of processed nodes

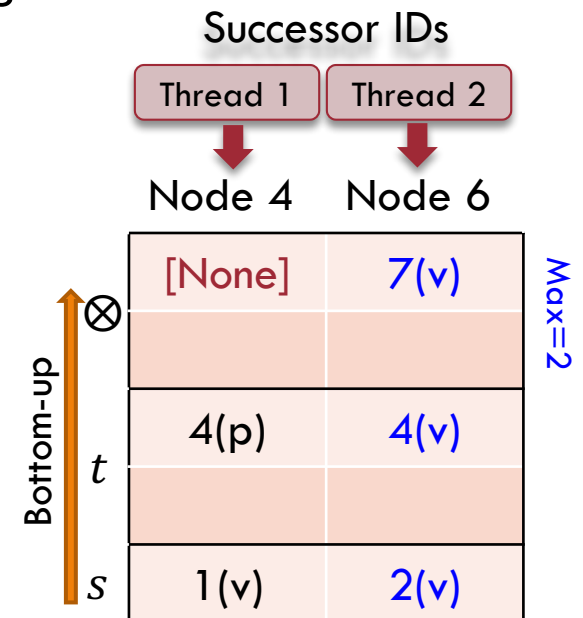




- Code-partitioning between the CPU and GPU according to the memory constraint
 - ▣ GPU lazily evaluates the Kronecker algebra operations
 - ▣ CPU maintains all computed thread interleavings

Lazy Evaluation on a GPU

23

- Execute lazy evaluation in 2 loops
 1. calculate node ID of all tree levels (top-down)
 2. retrieve successors (bottom-up)
- Intermediate results are stored in buffer
 - ▣ the size is pre-calculated from **max.** number of possible successors
 - ▣ align to maximize memory coalescing access

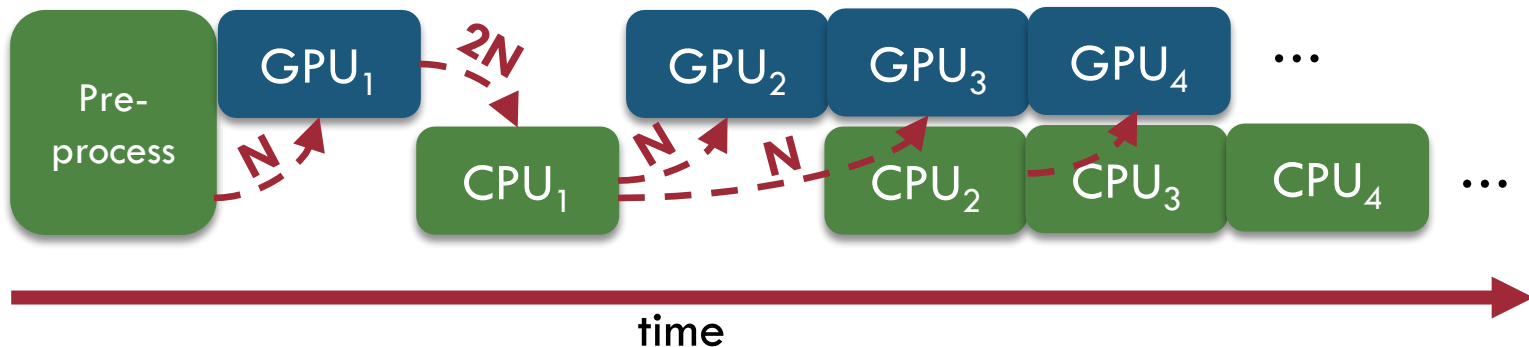


Pipelined Execution

24



- GPU processes N nodes per iteration and may discover $> N$ new successors
- The first N new successors are computed in the next iteration
- The remaining successors are computed in the following iteration
- The GPU computation can overlap with the CPU computation of the previous iteration




Experimental Setups

25

Experiments:

1. Dijkstra's Dining Philosophers
2. Linux kernel thread simulations
3. Railway system





Hardware specifications:

Platform Name	GTX 970	GTX 680	GT 750M	Xeon E5
CPU model	Intel i7-6700	Intel i7-3770k	Intel i7-4850HQ	Xeon E5-2697
CPU frequency	3.4 GHz	3.5 GHz	2.3 GHz	1.8 GHz
GPU model	NVIDIA GTX 970	NVIDIA GTX 680	NVIDIA GT 750M	N/A
GPU core frequency	1329 MHz	1006 MHz	822 MHz	
No. of GPU cores	1664	1536	384	
Compute Capability	5.2	3.0	3.5	

Experimental Results: CPU only vs. CPU+GPU

26





		3 kernel threads	5 kernel threads
GTX 970	KA-8	154 s	68024 s
	KA-G	0.30 s	14.64 s
GTX 680	KA-8	181 s	77107 s
	KA-G	0.29 s	14.14 s
GT 750M	KA-8	316 s	144321 s
	KA-G	0.52 s	69.10 s

- 12 Philosophers generate 1.6 million nodes
- 13 Philosophers generate 6.5 million nodes
- KA-G achieves up to **5453x** speedup over the fastest multi-threaded CPU implementation

Experimental Results: SPIN

27

Dining philosophers



		GTX 970	GTX 680	GT 750M
Railway	SPIN-8	1.64 s	1.68 s	2.05 s
	KA-8	4.15 s	4.89 s	5.86 s
	KA-G	0.05 s	0.05 s	0.07 s

- KA-G is faster than single-threaded SPIN
- Multi-threaded SPIN is unable to handle > 14 philosophers
- KA consumes less CPU memory
 - ▣ can handle larger problems

Conclusions

28

- Two-step lazy evaluation scheme to mitigate the state explosion problem
- Multicore CPU implementation
- Multicore CPU + GPU implementation
 - ▣ GPU: conducts lazy evaluation
 - ▣ CPU: maintains the computed thread interleavings
- The CPU+GPU implementation is up to **5453x** faster than our multicore CPU implementation
- Consumes up to **4.8x** less memory than SPIN-1 and **8.1x** less memory than SPIN-8

Thank you