

# Parallel Interval Newton Method on CUDA

Philip-Daniel Beck and Marco Nehmeier

Institute of Computer Science  
University of Würzburg  
Germany

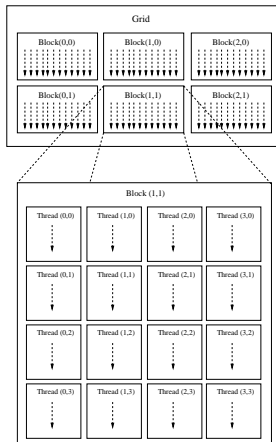
PARA 2012

- 1 CUDA
- 2 Interval Newton Method
- 3 Implementation on CUDA
- 4 Experimental Results
- 5 Summary

- 1 **CUDA**
- 2 Interval Newton Method
- 3 Implementation on CUDA
- 4 Experimental Results
- 5 Summary

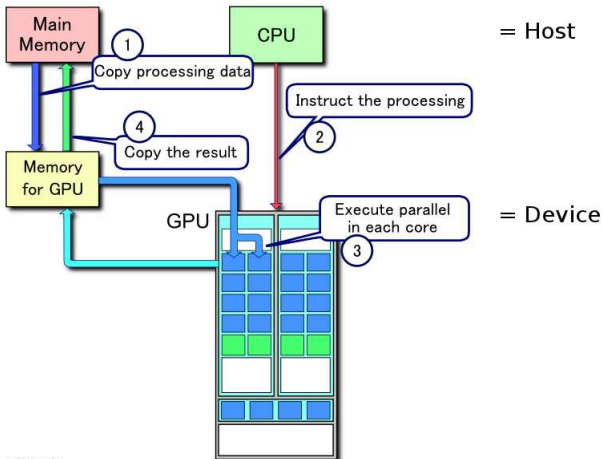
- GPU architecture for General Purpose Computing
- PCI Express add-in boards
- Highly parallel
  - Several streaming multiprocessors (SM)
    - Each consisting of several cores
- Compute capability = classification
- Usable like a many-core CPU
- CUDA C

- Grid = group of thread blocks
- Thread block = group of threads running on one SM
- Warp = 32 threads of a thread block
  - SIMT (Single Instruction Multiple Threads)
  - Fundamental unit



Memory	Access	Scope	Lifetime
Register	R/W	1 thread	Thread
Local	R/W	1 thread	Thread
Shared	R/W	All threads in block	Block
Global	R/W	All threads + host	Host allocation
Constant	R	All threads + host	Host allocation
Texture	R	All threads + host	Host allocation

# CUDA Processing flow



Author: Tosaka  
License: Creative Commons Attribution 3.0 Unported

- 1 CUDA
- 2 Interval Newton Method**
- 3 Implementation on CUDA
- 4 Experimental Results
- 5 Summary



$$N(X_n) := m(X_n) - \frac{f(m(X_n))}{F'(X_n)}$$

$$X_{n+1} := X_n \cap N(X_n)$$

- Idea: Use all tangents
- Reduce diameter
  - Use bisection otherwise
- Nested intervals
  - Find all zeros

1 Function  $f$

2 Interval  $X_n$

3  $0 \in F(X_n)$  ?

4  $N(X_n) := m(X_n) - \frac{f(m(X_n))}{F'(X_n)}$

5  $X_{n+1} := X_n \cap N(X_n)$


$X_{n+1} \neq \emptyset$  ?

$\Rightarrow$  Rec. call with  $X_{n+1}$

$X_{n+1} = X_n$  ?

$\Rightarrow$  Bisection of  $X_{n+1}$


$\Rightarrow$  Two rec. calls



- 1 Function  $f$
- 2 Interval  $X_n$
- 3  $0 \in F(X_n)$  ?
- 4  $N(X_n) := m(X_n) - \frac{f(m(X_n))}{F'(X_n)}$
- 5  $X_{n+1} := X_n \cap N(X_n)$

$X_{n+1} \neq \emptyset$  ?  
 $\Rightarrow$  Rec. call with  $X_{n+1}$


$X_{n+1} = X_n$  ?  
 $\Rightarrow$  Bisection of  $X_{n+1}$   
 $\Rightarrow$  Two rec. calls



- 1 Function  $f$
- 2 Interval  $X_n$
- 3  $0 \in F(X_n)$  ?
- 4  $N(X_n) := m(X_n) - \frac{f(m(X_n))}{F'(X_n)}$
- 5  $X_{n+1} := X_n \cap N(X_n)$

$X_{n+1} \neq \emptyset$  ?  
 $\Rightarrow$  Rec. call with  $X_{n+1}$


$X_{n+1} = X_n$  ?  
 $\Rightarrow$  Bisection of  $X_{n+1}$   
 $\Rightarrow$  Two rec. calls



- 1 Function  $f$
- 2 Interval  $X_n$
- 3  $0 \in F(X_n)$  ?
- 4  $N(X_n) := m(X_n) - \frac{f(m(X_n))}{F'(X_n)}$
- 5  $X_{n+1} := X_n \cap N(X_n)$

$X_{n+1} \neq \emptyset$  ?  
 $\Rightarrow$  Rec. call with  $X_{n+1}$


$X_{n+1} = X_n$  ?  
 $\Rightarrow$  Bisection of  $X_{n+1}$   
 $\Rightarrow$  Two rec. calls



- 1 Function  $f$
- 2 Interval  $X_n$
- 3  $0 \in F(X_n)$  ?
- 4  $N(X_n) := m(X_n) - \frac{f(m(X_n))}{F'(X_n)}$
- 5  $X_{n+1} := X_n \cap N(X_n)$

$X_{n+1} \neq \emptyset$  ?  
 $\Rightarrow$  Rec. call with  $X_{n+1}$


$X_{n+1} = X_n$  ?  
 $\Rightarrow$  Bisection of  $X_{n+1}$   
 $\Rightarrow$  Two rec. calls



- 1 Function  $f$
- 2 Interval  $X_n$
- 3  $0 \in F(X_n)$  ?
- 4  $N(X_n) := m(X_n) - \frac{f(m(X_n))}{F'(X_n)}$
- 5  $X_{n+1} := X_n \cap N(X_n)$

$X_{n+1} \neq \emptyset$  ?  
 $\Rightarrow$  Rec. call with  $X_{n+1}$


$X_{n+1} = X_n$  ?  
 $\Rightarrow$  Bisection of  $X_{n+1}$   
 $\Rightarrow$  Two rec. calls



- 1 Function  $f$
- 2 Interval  $X_n$
- 3  $0 \in F(X_n)$  ?
- 4  $N(X_n) := m(X_n) - \frac{f(m(X_n))}{F'(X_n)}$
- 5  $X_{n+1} := X_n \cap N(X_n)$

$X_{n+1} \neq \emptyset$  ?  
 $\Rightarrow$  Rec. call with  $X_{n+1}$

$X_{n+1} = X_n$  ?  
 $\Rightarrow$  Bisection of  $X_{n+1}$   
 $\Rightarrow$  Two rec. calls






- 1 Function  $f$
- 2 Interval  $X_n$
- 3  $0 \in F(X_n)$  ?
- 4  $N(X_n) := m(X_n) - \frac{f(m(X_n))}{F'(X_n)}$
- 5  $X_{n+1} := X_n \cap N(X_n)$

$X_{n+1} \neq \emptyset$  ?  
 $\Rightarrow$  Rec. call with  $X_{n+1}$


$X_{n+1} = X_n$  ?  
 $\Rightarrow$  Bisection of  $X_{n+1}$   
 $\Rightarrow$  Two rec. calls



Non monotonic ?

⇒ Extended interval arithmetic


⇒ Two recursive calls



Non monotonic ?


⇒ Extended interval arithmetic

⇒ Two recursive calls



Non monotonic ?

- ⇒ Extended interval arithmetic
- ⇒ Two recursive calls



- 1 CUDA
- 2 Interval Newton Method
- 3 Implementation on CUDA**
- 4 Experimental Results
- 5 Summary

- Sequential method
- Static load balancing not advisable
  - Bisection of the work for the cases
    - $X_{n+1} = X_n$
    - $0 \in F'(X_n)$
  - Idleness for the cases
    - $0 \notin F(X_n)$
    - $X_{n+1} = \emptyset$

⇒ We need dynamic load balancing!

- Sequential method
- Static load balancing not advisable
  - Bisection of the work for the cases
    - $X_{n+1} = X_n$
    - $0 \in F'(X_n)$
  - Idleness for the cases
    - $0 \notin F(X_n)$
    - $X_{n+1} = \emptyset$

⇒ We need dynamic load balancing!

- Blocking Queue
- Task Stealing
- Distributed Stacks
- Static Task List

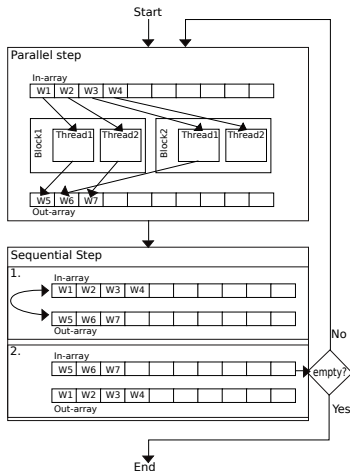


- Global Queue
- Mutual exclusion
- Persistent threads

- Lock free
- Block local queue
- Steal tasks from other blocks in case of an empty queue
- Persistent threads

- Lock free
- Block local stack
- Load balancing only between threads of a block
- Static distribution of the workload at the beginning
- Persistent threads

- Lock free
- Two Arrays
  - In-array (read)
  - Out- array (write)
- Swap arrays after each iteration
  - Terminate kernel
  - Initialize new kernel with swapped arrays
    - Number of blocks depends on array size



- 1 CUDA
- 2 Interval Newton Method
- 3 Implementation on CUDA
- 4 Experimental Results**
- 5 Summary

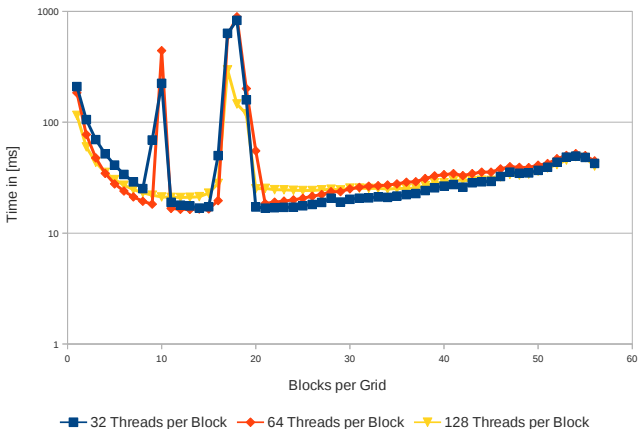
- Intel Xeon E5504 2.00 GHz
- 8 GB RAM
- Debian 64 Bit Linux
- NVIDIA Tesla C2070
  - Compute capability 2.0
  - 14 streaming multiprocessors
    - 32 cores
  - 6 GB RAM
  - 515 GigaFLOPS (double)
  - 1.03 TFLOPS (single)
- NVCC 3.2
- GNU gcc 4.4.5
- Boost 1.46.1
- filib++ 3.0.1

- NVCC `-ptxas-options=-v`

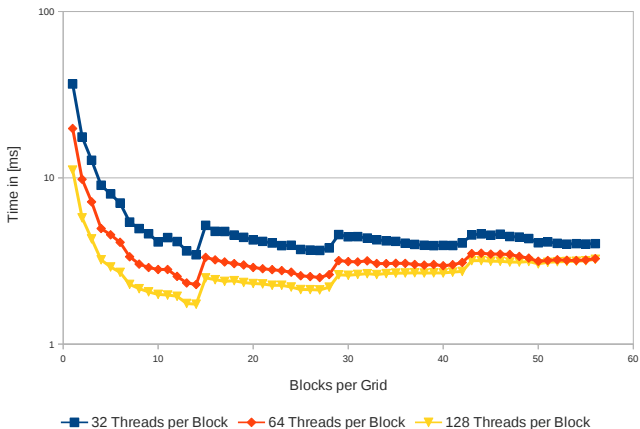
Method	Register/Thread	Shared Memory/Block
BlockingQueue	63	8240 Byte
TaskStealing	63	8240 Byte
DistributedStacks	63	32784 Byte
StaticTaskList	59	0 Byte

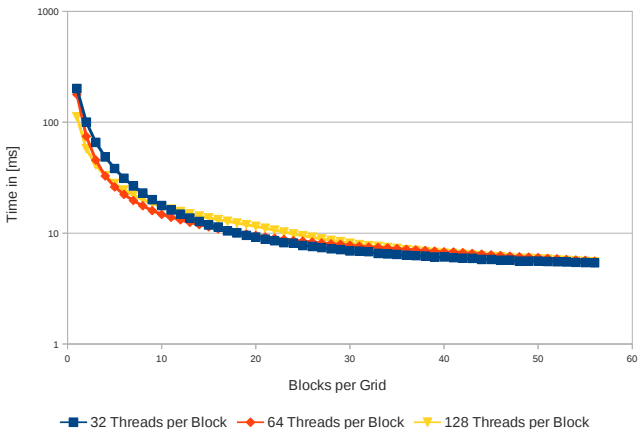
- Compute Capability 2.0

- 32768 Registers
- 48 KB Shared Memory










Method	Blocks per grid	Threads per block
BlockingQueue	14	64
TaskStealing	28	64
DistributedStacks	14	128
StaticTaskList	-	32


$\sinh x$

$X_0 = [-1, 1], \epsilon = 10^{-12}$




$$\sin x - \frac{x}{100}$$

$$X_0 = [0, 100], \epsilon = 10^{-12}$$



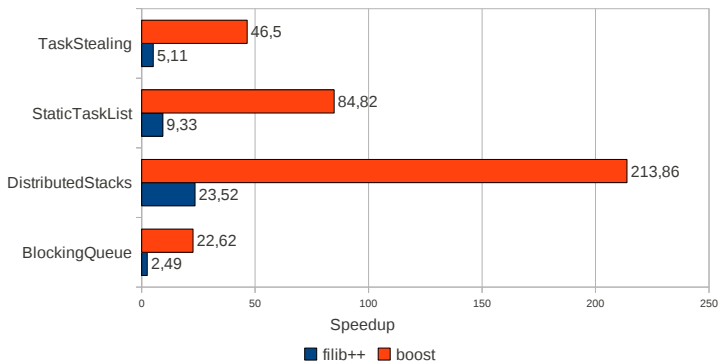
$$\sin x - \frac{x}{10000}$$

$$X_0 = [0, 10000], \epsilon = 10^{-12}$$




$$\sin x - \frac{x}{10000}$$

$$X_0 = [0, 10000], \epsilon = 10^{-12}$$



$$\sin \frac{1}{x}$$


$$X_0 = [0.00001, 15], \epsilon = 10^{-12}$$






$$\sin \frac{1}{x}$$

$$X_0 = [0.00001, 15], \epsilon = 10^{-12}$$




$$(3x^3 - 5x + 2) \cdot \sin^2 x + (x^3 + 5 \cdot x) \cdot \sin x - 2x^2 - x - 2$$

$$X_0 = [-10, 10], \epsilon = 10^{-12}$$




$$x^{14} - 539.25 \cdot x^{12} + 60033.8 \cdot x^{10} - 1.77574e^6 \cdot x^8 + 1.70316e^7 \cdot x^6 - 5.50378e^7 \cdot x^4 + 4.87225e^7 \cdot x^2 - 9.0e^6$$

$$X_0 = [-20, 20], \epsilon = 10^{-12}$$




$$(3x^3 - 5x + 2) \cdot \sin^2 x + (x^3 + 5 \cdot x) \cdot \sin x - 2x^2 - x - 2$$

$$X_0 = [-10, 10], \epsilon = 10^{-6}$$




$\sinh x$

$X_0 = [-1, 1], \epsilon = 10^{-12}$




$$\sin \frac{1}{x}$$

$$X_0 = [0.00001, 15], \epsilon = 10^{-12}$$



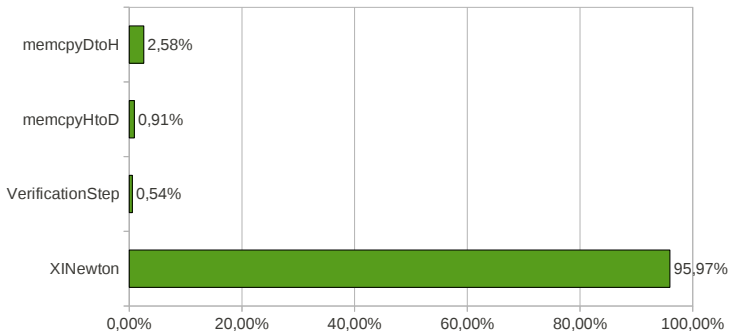
$\sinh x$ 

$$X_0 = [-1, 1], \epsilon = 10^{-12}$$



$$\sin \frac{1}{x}$$


$$X_0 = [0.00001, 15], \epsilon = 10^{-12}$$






$$\sin x - \frac{x}{100}$$

$$X_0 = [0, 100], \epsilon = 10^{-12}$$



$$\sin \frac{1}{x}$$

$$X_0 = [0.00001, 15], \epsilon = 10^{-12}$$



- 1 CUDA
- 2 Interval Newton Method
- 3 Implementation on CUDA
- 4 Experimental Results
- 5 Summary**

- Parallel interval Newton method
- Many roots
- Global memory not advisable
- Static Task List
  - Slow for simple functions
- Distributed Stacks

# Questions ?

Marco Nehmeier

[nehmeier@informatik.uni-wuerzburg.de](mailto:nehmeier@informatik.uni-wuerzburg.de)