



ISC 2011

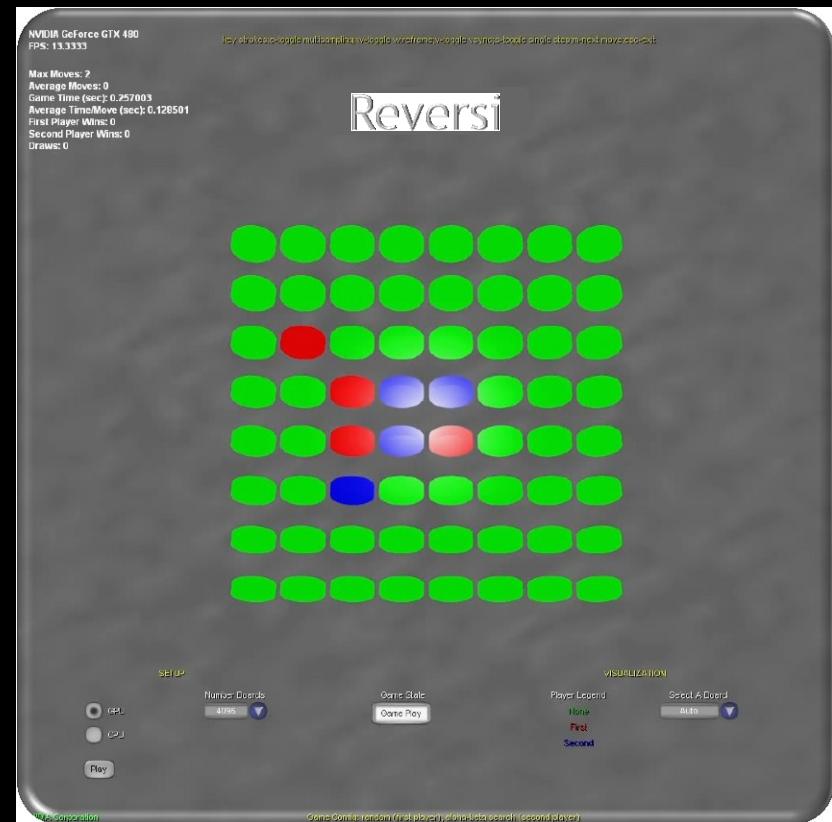
Producer-Consumer Model for Massively Parallel Zero-Sum Games on the GPU

Avi Bleiweiss
NVIDIA Corporation

Zero-Sum

- One's gain \Rightarrow other's loss
- Perfect info
- Multi player game
 - Simple and involved

Matching Pennies	Head	Tail
Head	1, -1	-1, 1
Tail	-1, 1	1, -1



Motivation

- Play as you go
 - Thousands players
- Game cloud
 - GPU computing
- Mobile computer
 - Commodity

NVIDIA Tesla



NVIDIA Tegra

Problem

Game

- Two player
- Maximize look ahead
- Rapid node expansion
 - 10^{25} for 4x4x4 Tic-Tac-Toe
 - 10^{120} for Chess

Search

- Efficient parallel
 - Space split
 - Statistical simulation
 - Simultaneous matches
- Many thousands threads

Parallelism

Principal Variation Split

- Strongly ordered tree
- Synchronization bound
- Load imbalance

Young Brothers Wait Concept

- Parallel at any node
- Processor owns node
- Up to 1024 processors

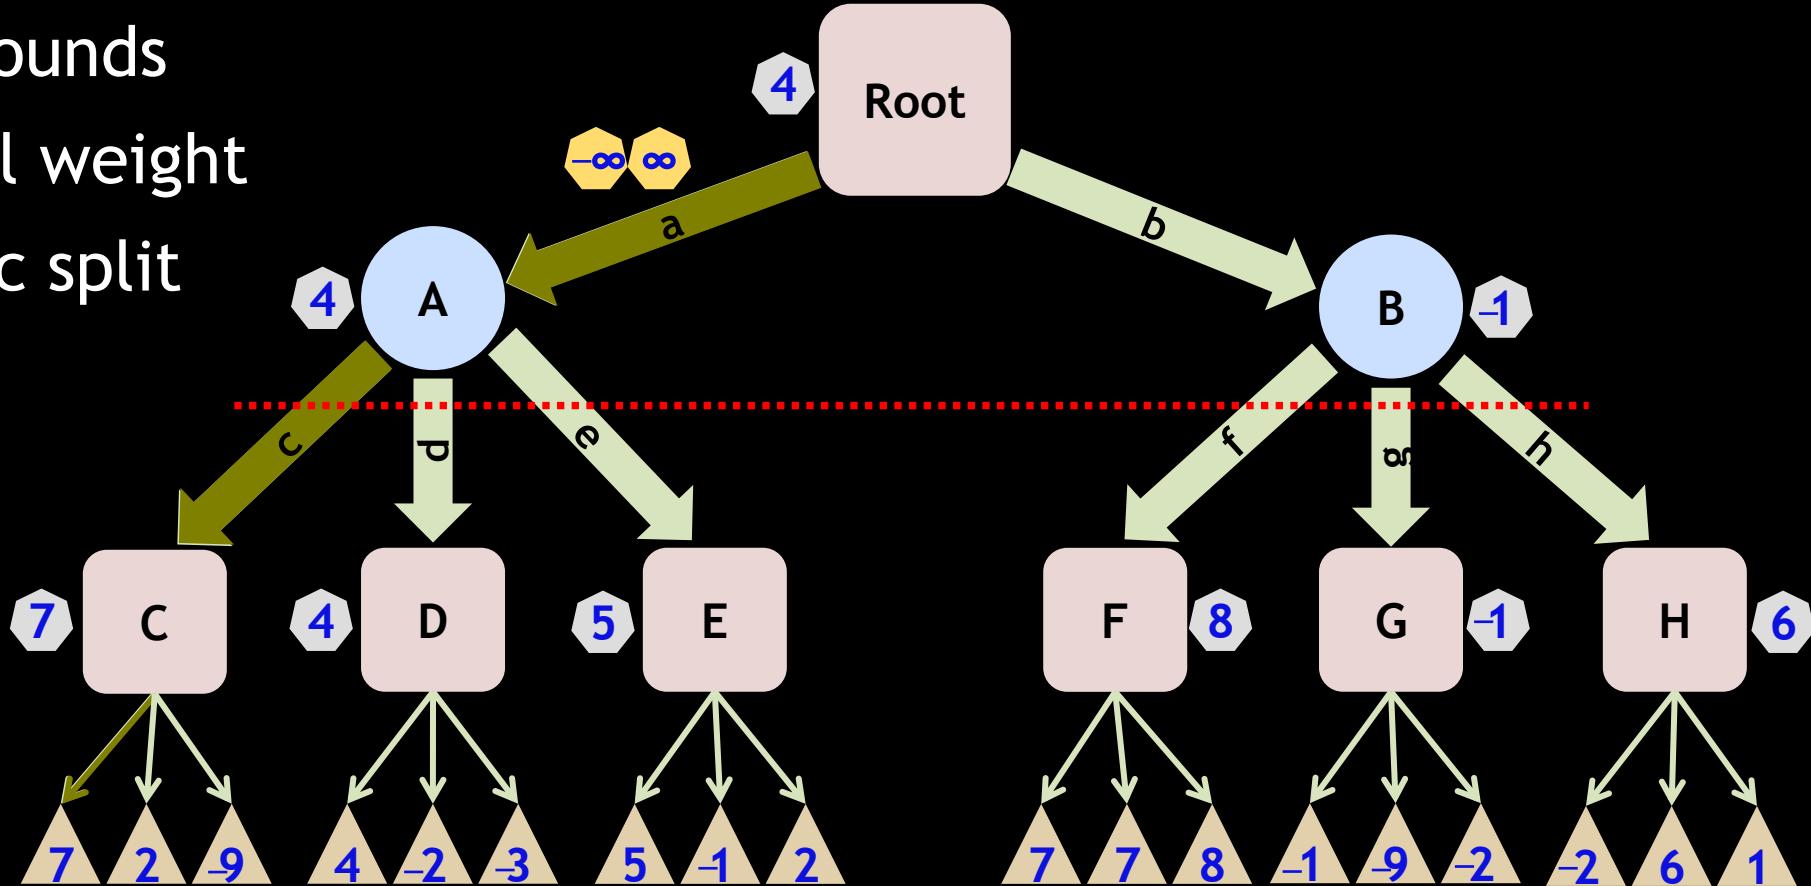
Dynamic Tree Splitting

- Processors share node
- Global job list
- Reasonable speedup

No massive parallel solution in shared memory settings!

Cut Nodes

- PV Bounds
- Equal weight
- Static split



Heuristic Search

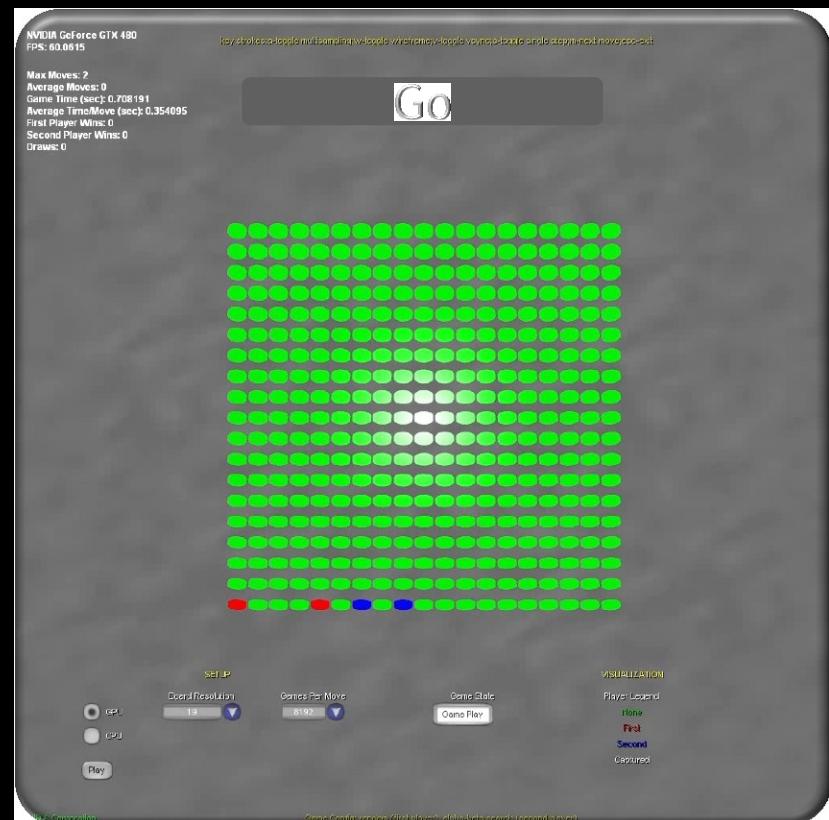
```
AlphaBeta( $v, \alpha_l, \beta_l, \alpha_g, \beta_g$ )
 $\alpha_l \leftarrow \max(\alpha_l, \alpha_g)$ 
 $\beta_l \leftarrow \min(\beta_l, \beta_g)$ 
if ( $\alpha_l > \beta_l$ ) return rank
if ( $s(v) = \emptyset$ ) return  $f(v)$ 
foreach ( $v_j \ni s(v)$ ) do
    value  $\leftarrow$  AlphaBeta( $v_j, \alpha_l, \beta_l, \alpha_g, \beta_g$ )
    if ( $p(v) == First$ ) then
        rank  $\leftarrow \max(rank, value)$ 
         $\alpha_l \leftarrow \max(\alpha_l, rank)$ 
        if ( $\alpha_l > \alpha_g$ ) then  $\alpha_g \xleftarrow{atomic} \alpha_l$ 
    else
        rank  $\leftarrow \min(rank, value)$ 
         $\beta_l \leftarrow \min(\beta_l, rank)$ 
        if ( $\beta_l < \beta_g$ ) then  $\beta_g \xleftarrow{atomic} \beta_l$ 
return rank
```

Statistical Simulation

- Monte Carlo method
 - Annealing process
- k -simulation move error

$$\Delta f(v) \sim \frac{1}{\sqrt{k}}$$

- Many thousands games
 - for 19x19 Go game



Challenges

Deep recursion, limited stack

Divergent, irregular threads

Dynamic parallelism

Low arithmetic intensity

Implementation

- Kernel for each
 - Alpha-Beta, Monte Carlo
- Board C++ class
 - Rules specific
- Games

Heuristic
Search

3D Tic-Tac-Toe

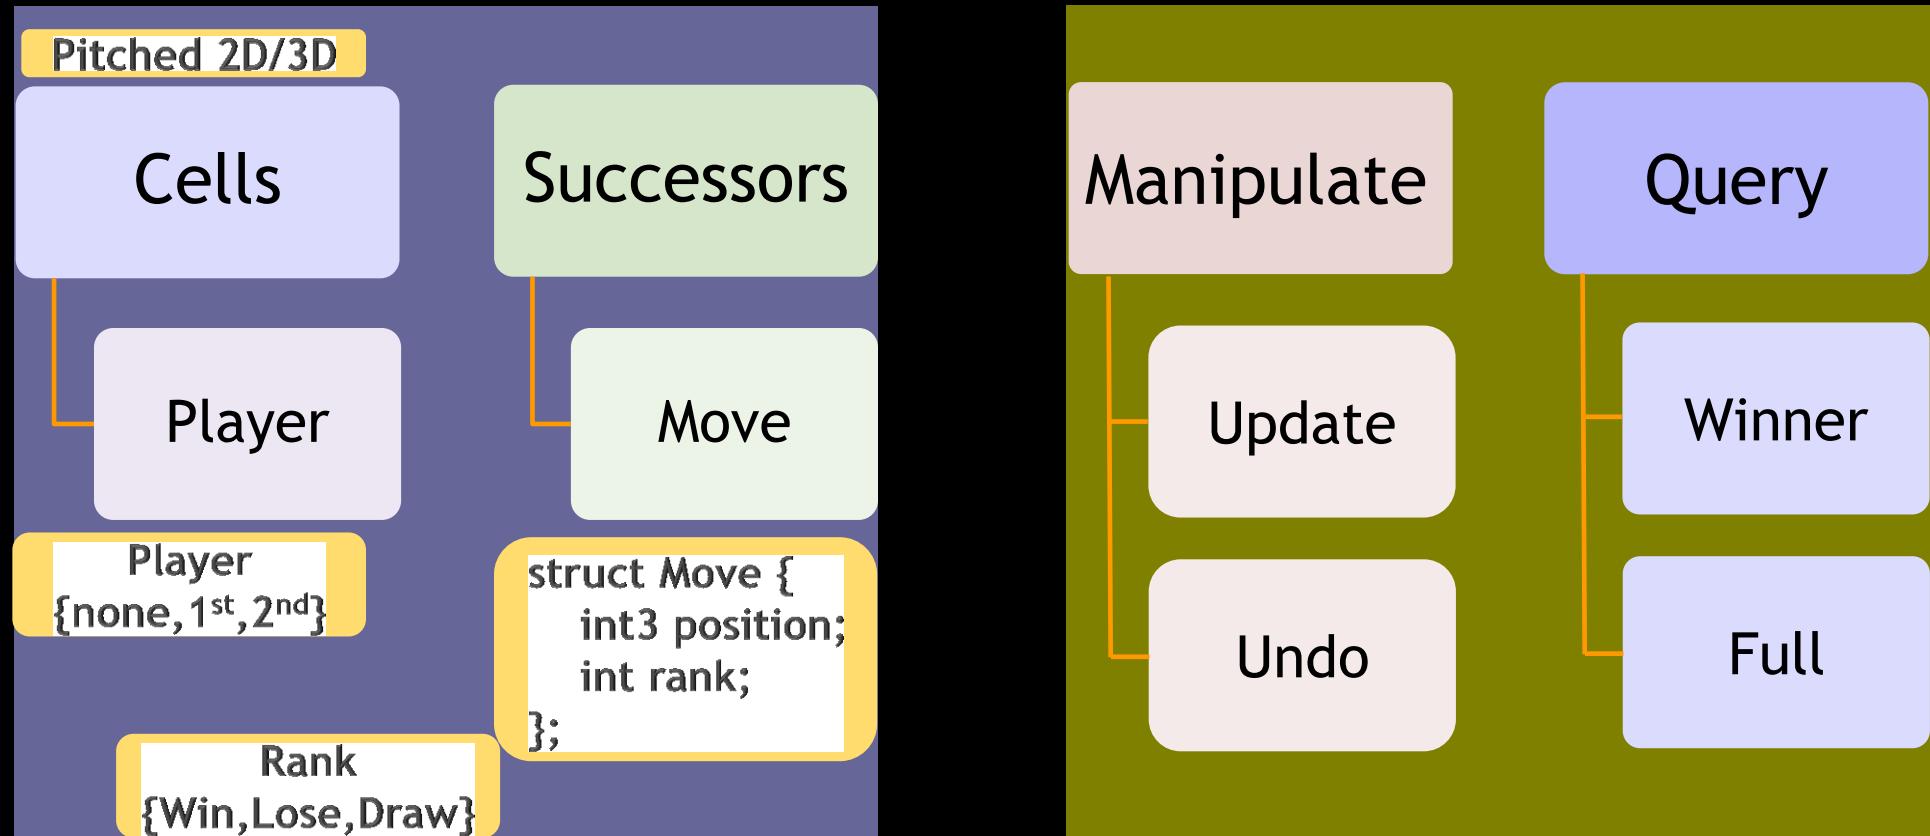
Connect-4

Reversi

Monte Carlo
Simulation

Go

State Abstraction



Stack

- Recursion depth >1000
- Greedy allocation

*StackBytes/Thread * SM # * MaxWarps/SM * Threads/Warp*

- Hybrid design

Runtime/Compiler

User

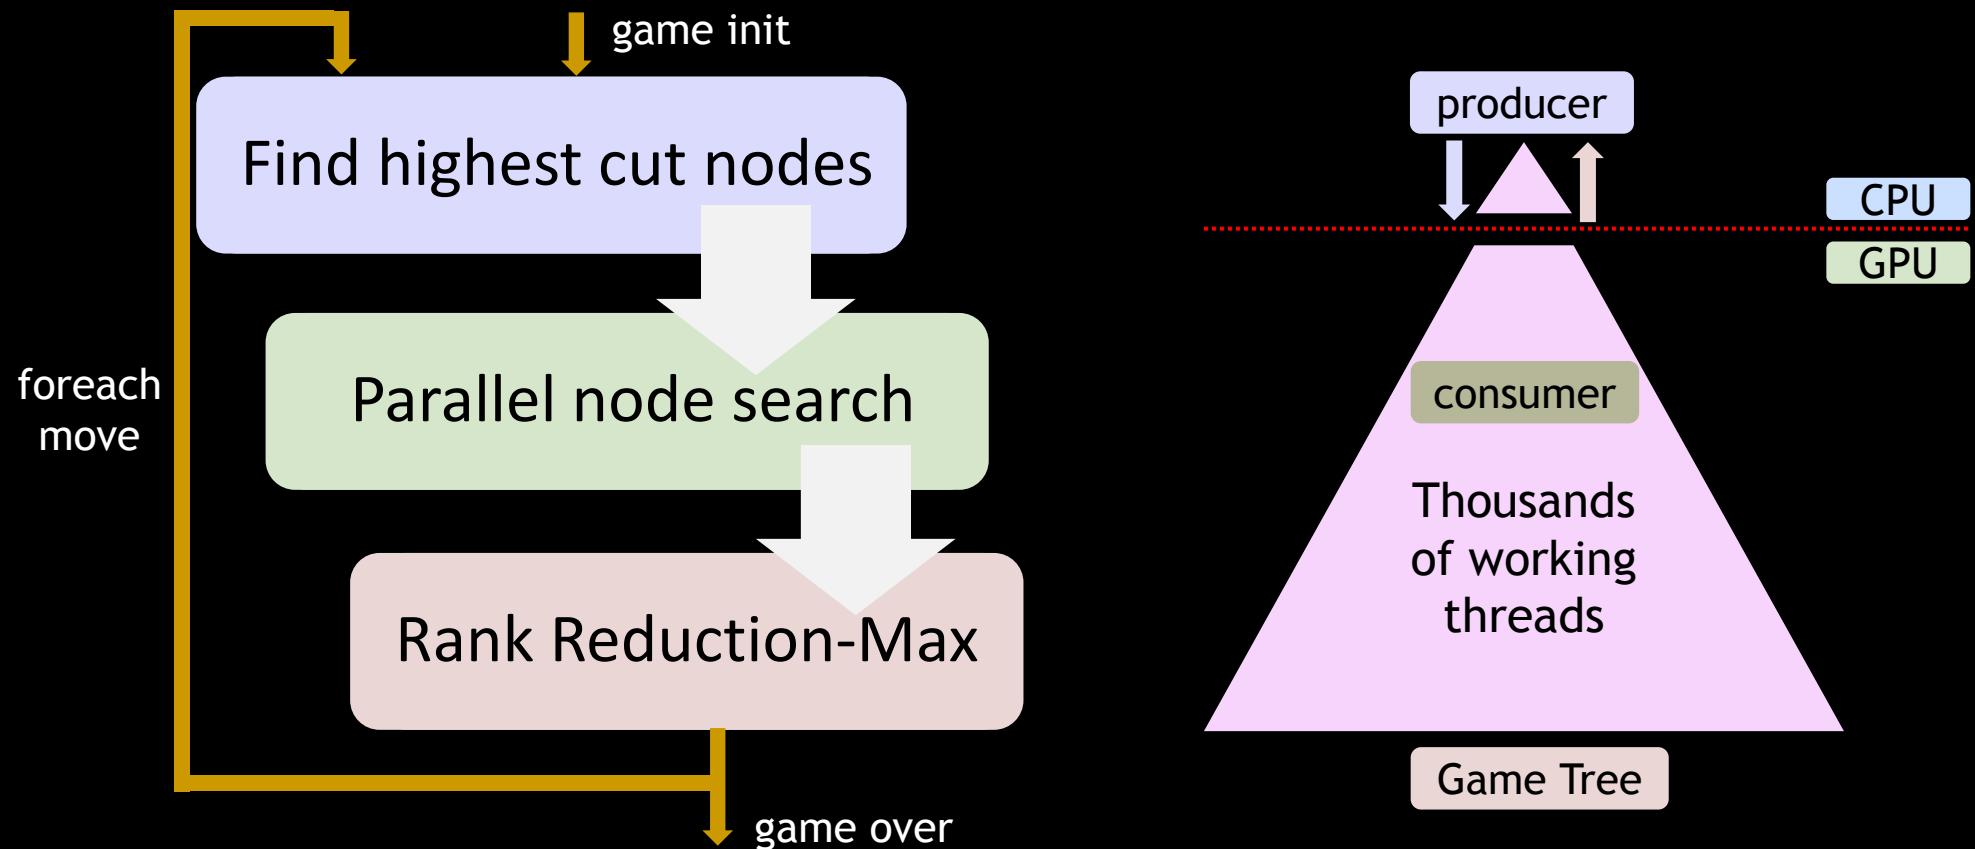
Local
Memory

- Local variables
- Function parameters

- Successors

Global
Memory

Producer-Consumer



Shared αβ

Kernel global scope

Check local αβ

```
__device__ int *galpha, *gbeta;  
  
__device__ void  
resolve(int& alpha, int& beta)  
{  
    if(alpha <= *galpha) alpha = *galpha;  
    else atomicMax(galpha, alpha);  
  
    if(beta >= *gbeta) beta = *gbeta;  
    else atomicMin(gbeta, beta);  
}
```

Global atomic update

Limitations

- Stack allocation
 - Bounds parallelism
- Static split constraints

Depth	Tic-Tac-Toe Threads		
	3x3x3	4x4x4	5x5x5
1	650	3906	15252
2	15600	238266	1860744

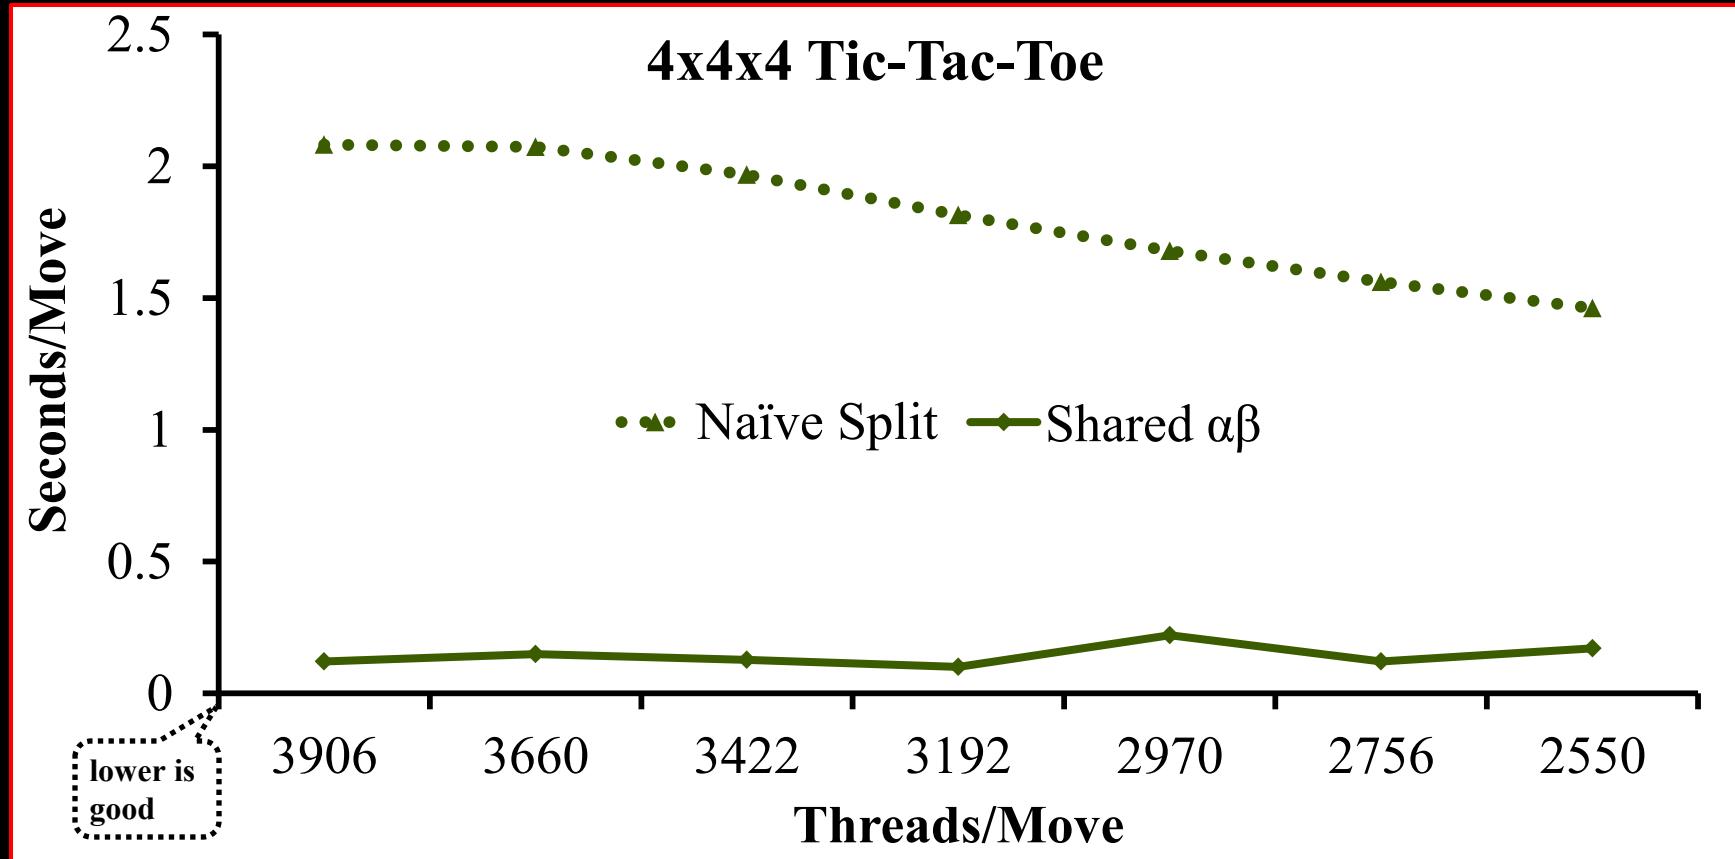
Methodology

- CUDA Toolkit 3.1, Windows
- Single processor

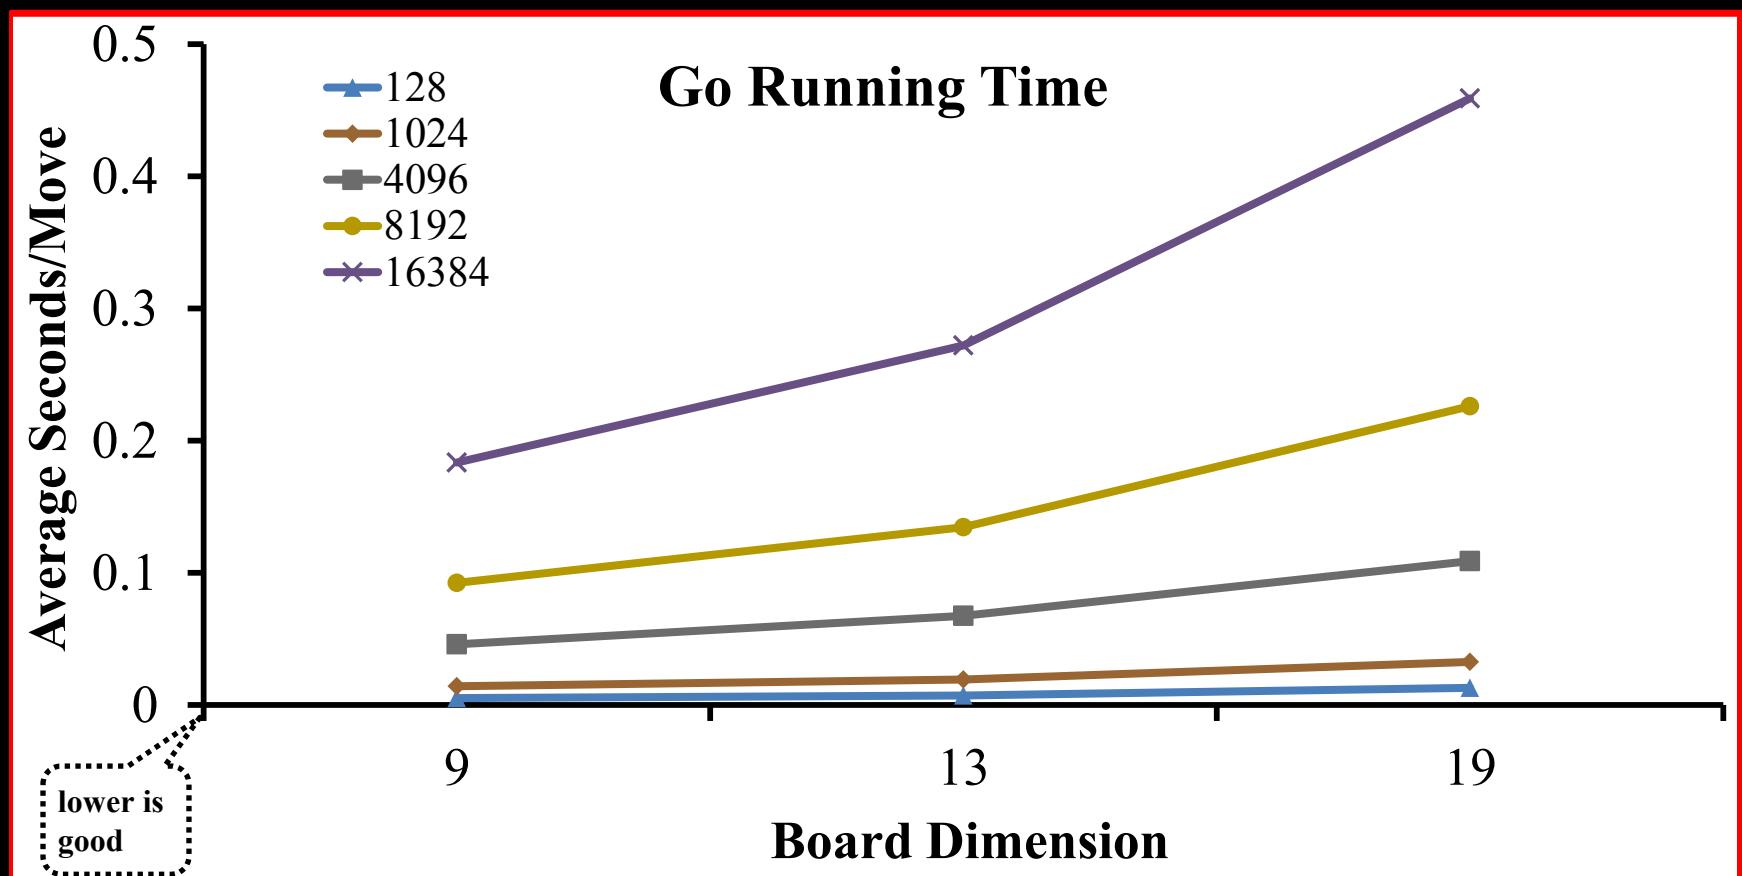
GPU	SMs	Warps/SM	Clocks(MHz)	L1/Shared (KB)	L2(KB)
GTX480	15	2	723/1446/1796	48/16	640

CPU	Cores	Clocks(MHz)	L1/L2 (KB)
I7-940	8	2942/(3*1066)	32/8192

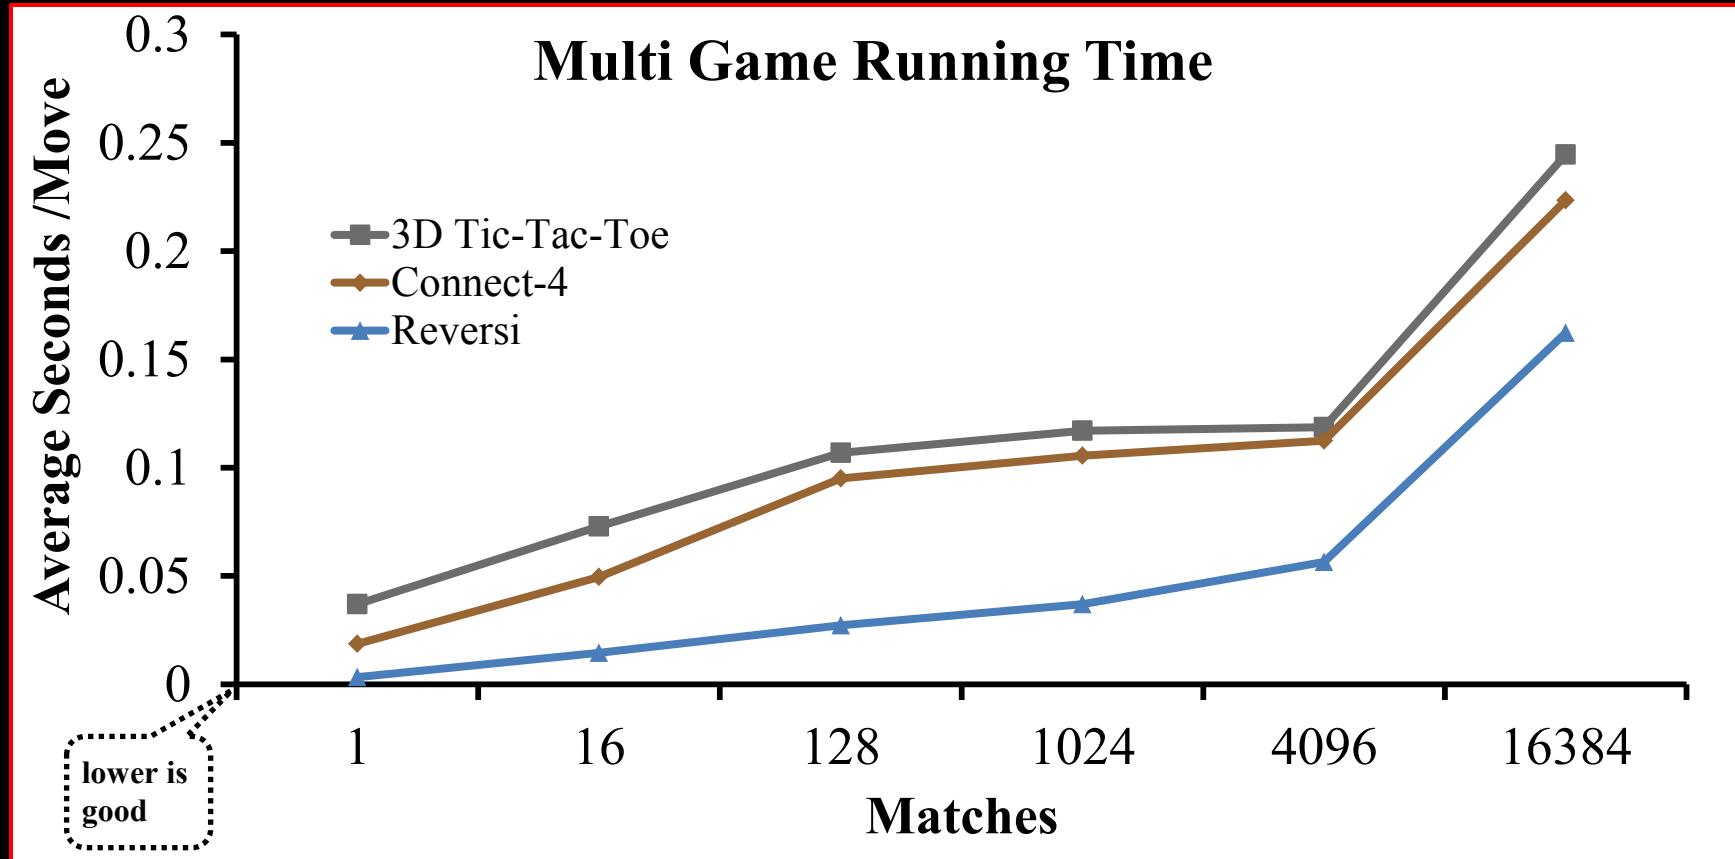
Space Split



Monte Carlo

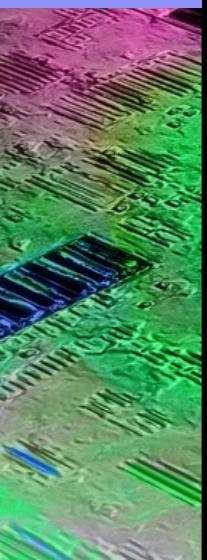


Simultaneous Matches



Future Work

- Data packing optimization
- Per-SM transposition table
- Predictable node ranks
- Dynamic space split

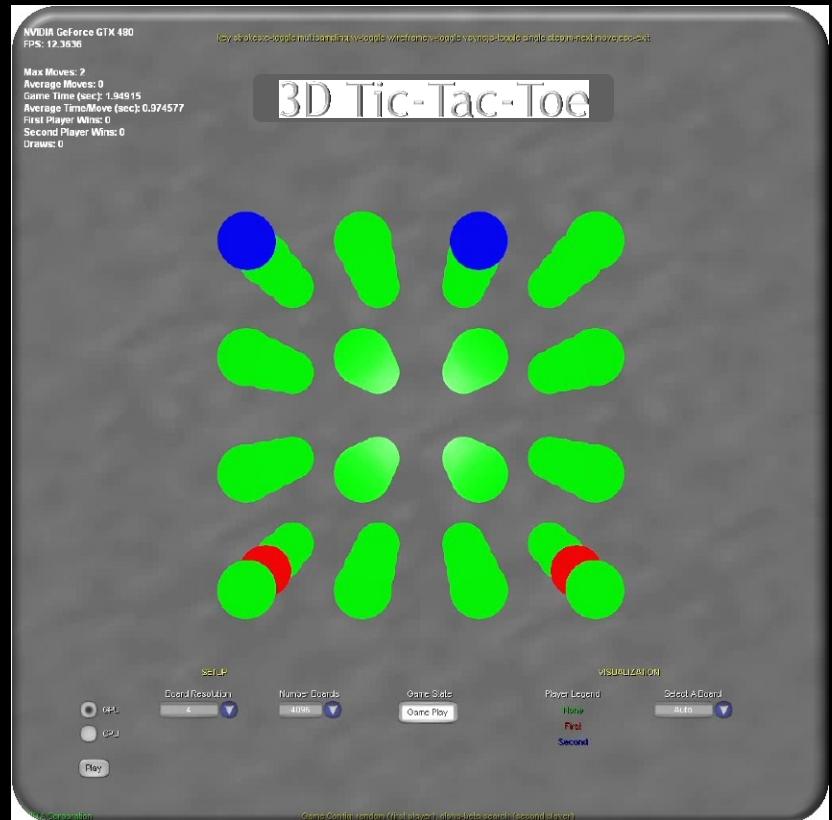


GPU Performance

Metric	Game	Dimension	Speedup
Shared αβ vs. Naïve Split	3D Tic-Tac-Toe	4x4x4	13.37X mean
Monte Carlo vs. CPU	Go	19x19	121.64X @16K

Summary

- Efficient GPU based
 - Heuristic search
 - Statistical Simulation
- Economical solution
 - Mobile-Cloud



Thank You!



Questions?

Info

- Base
 - <http://developer.nvidia.com>
- GPU AI for Board Games
 - [Technology Preview](#)
- Toolkit
 - [CUDA Zone](#)
- Debugger
 - [Parallel Nsight](#)