

Parallel A* pathfinding algorithm

Giuseppe Accaputo
Pascal Iselin

December 16, 2013

Overview

- ▶ Goal
- ▶ Literature Review
- ▶ Approaches
- ▶ Results

Our Goals were

- ▶ Implement a correct parallel A* algorithm
- ▶ Make it faster than the serial version

Our Goals were


- ▶ Implement a correct parallel A* algorithm ✓
- ▶ Make it faster than the serial version

Our Goals were

- ▶ Implement a correct parallel A* algorithm ✓
- ▶ Make it faster than the serial version ✗

Best first heuristic search

$$f(n) = \underbrace{g(n)}_{\text{exact cost}} + \underbrace{h(n)}_{\text{estimated cost}}$$



Parallel A* in the literature

- ▶ **Shared Priority Queue:** Threads produce and consume simultaneously. Does not perform well!
[Cohen et al., 2010]
- ▶ **Bidirectional Search:** Run two searches. Does not scale!
[Rios, Luis Henrique Oliveira, and Luiz Chaimowicz. PNBA: A Parallel Bidirectional Heuristic Search Algorithm.]*
- ▶ **Sacrifice path quality for speed:** Converge towards other algorithms
[Sandy Brand, and Rafael Bidarra. Multicore scalable and efficient pathfinding with Parallel Ripple Search. Comp. Anim. Virtual Worlds 23.2 (2012)]
- ▶ **Clustering:** Too complicated
[Rafia, Inam. A Algorithm for Multicore Graphics Processors. (2010).]*

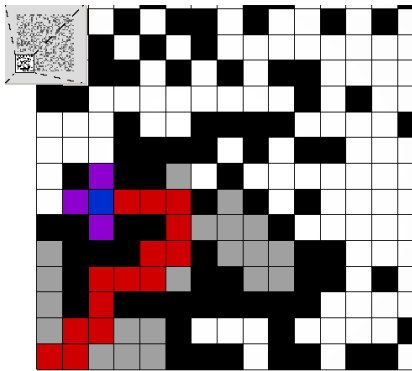
Ideas we implemented

- ▶ Concurrent Neighbor Expansion
- ▶ Shared Priority Queue
- ▶ Atomic ClosedFlags + Shared JobQueue

Same underlying datastructure for all the Implementations.
SquareLattice filled with Objects of Type MapNode

Concurrent Neighbor Expansion

Concurrent calculation of the neighbors in each step. This does not scale but it's easy!




Shared Priority Queue

- ▶ *concurrent_priority_queue* (CPQ) from Intel's Thread Building Blocks (TBB)
- ▶ CPQ does not allow rebalancing
- ▶ Recursive call of the search function with a counter
- ▶ One lock per node

```
1 atomic <size_t> num_threads;
2 task_group t_group;
3
4 t_group.run([&]{ parallel_search(); });
5
6 void parallel_search() {
7     //-----
8     // A* Magic
9     //-----
10    size_t n_threads = ++num_threads;
11    if(n_threads < max_threads) {
12        t_group.run([&]{ parallel_search(); });
13    } else {
14        --num_threads;
15    }
16 }
```

Atomic ClosedFlags + Shared JobQueue

- ▶ Run serial A* until we have enough open nodes
- ▶ Run a new A* in parallel on each of the open nodes
- ▶ Each thread has its own priority queue
- ▶ Threads communicate through a shared grid of closed flags
- ▶ Use atomic CAS to set the flags




Atomic ClosedFlags + Shared JobQueue


- ▶ Run serial A* until we have enough open nodes
- ▶ Run a new A* in parallel on each of the open nodes
- ▶ Each thread has it's own priority queue
- ▶ Threads communicate through a shared grid of closed flags
- ▶ Use atomic CAS to set the flags

- ▶ How to make sure threads don't just terminate? → Shared JobQueue
- ▶ How to guarantee the shortest path?

How to guarantee the shortest path?



How to guarantee the shortest path?




How to guarantee the shortest path?

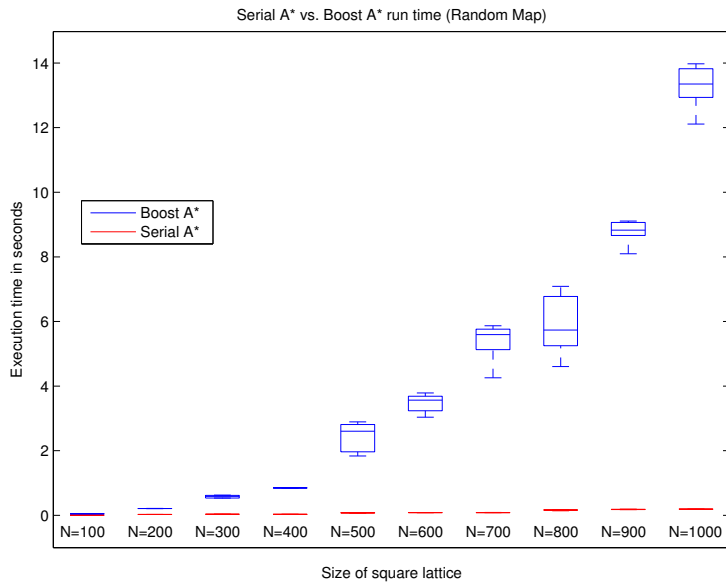


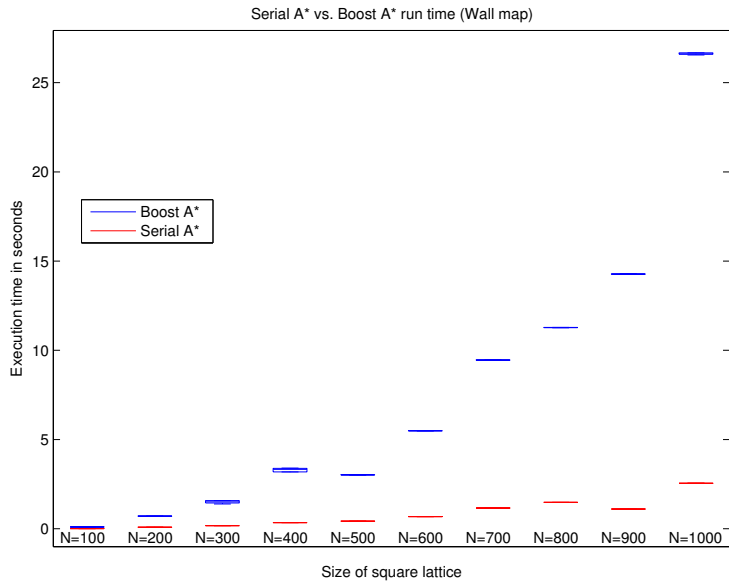
Just take the green path... **NO!** That would be fringe search!

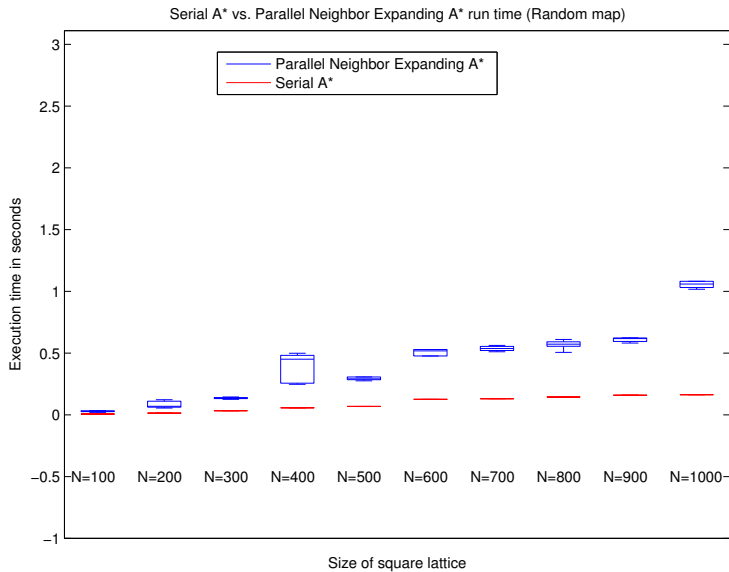
Test setup

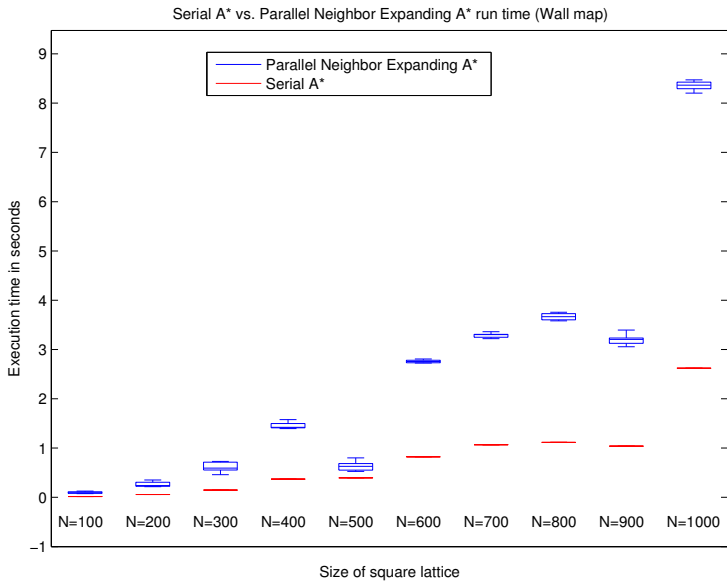
- ▶ Tested on Kanifushi
(32 Cores Intel Xeon E7-4830 @ 2.13Ghz)
- ▶ Compiled with GCC 4.7 and TBB 3.0
- ▶ 10 runs per test case
- ▶ 90% of random map was walkable




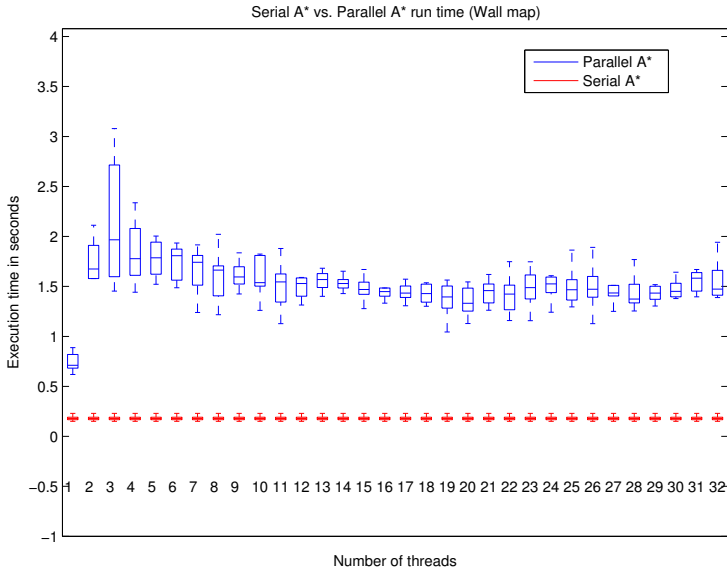


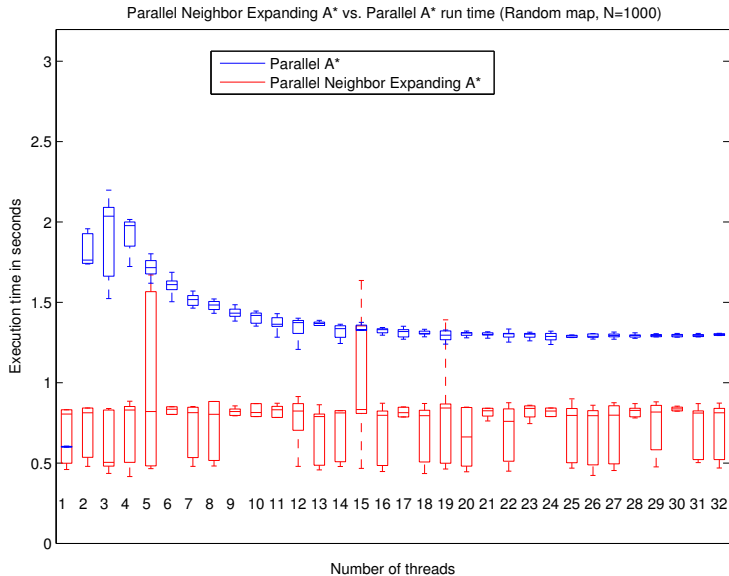


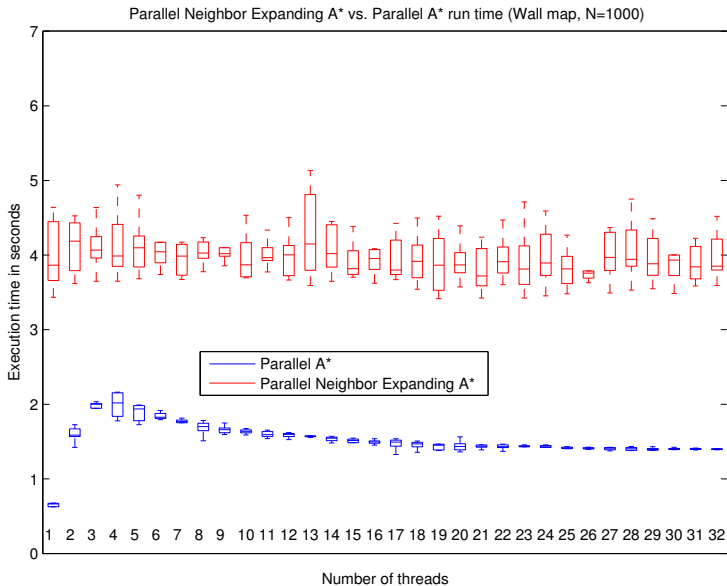












Conclusions

- ▶ We comply with the literature!
- ▶ One must sacrifice path quality for speed
- ▶ There are much better alternatives out there:
 - ▶ Ripple Search [*Brand et al., 2012*]
 - ▶ Fringe Search