

Envelope-based Approaches to Real-Time Heuristic Search

Kevin C. Gall, Bence Cserna, Wheeler Ruml
kcg245 at gmail.com, bence, ruml at cs.unh.edu

Motivation

- Interactive systems often must be controlled in real time requiring a real-time planner.
- Prevailing paradigm is ‘agent-centered’ which leads to ‘scrubbing’ in local minima.
- Real-time systems should be able to direct the agent toward the best state encountered in the search so far.

Envelope Search

- Maintain a single search envelope.
- Escape local minima without ‘scrubbing.’
- Core problem: How does the agent efficiently connect to the envelope frontier?

Contributions

- An analysis of existing envelope-based techniques: TBA* and its variants.
- A new envelope-based algorithm that addresses the flaws of TBA* by using an auxiliary search to guide agent to the frontier.
- Empirically demonstrate our algorithm achieves state-of-the-art performance in domains with dense state space graphs.

Time Bounded A* (IJCAI, 2009)

- Single A* search across planning iterations
- Agent is committed along the best discovered branch of the A* search tree.
- **TB Weighted A*** $TB(WA^*)$ and **TB Best First Search** $TB(BFS)$ - bias the search graph toward falling h values allowing the agent to quickly explore suboptimal branches.
- Agent’s path is always rooted in the start state.
- Irrational behavior stems from the agent’s strict adherence to a rigid search graph.

Intra-Envelope Search (I-ES)

- A new Envelope-based search algorithm that does not rely solely on the A* search tree.
- Periodically select a promising node on the envelope frontier as the agent’s target.
- Backward search connects s_{target} to s_{agent} guided by $h(s_{agent}, s_{target})$.
- Once the backward search reaches the agent, path to target is cached. Subsequent searches terminate when they reach the cached path.

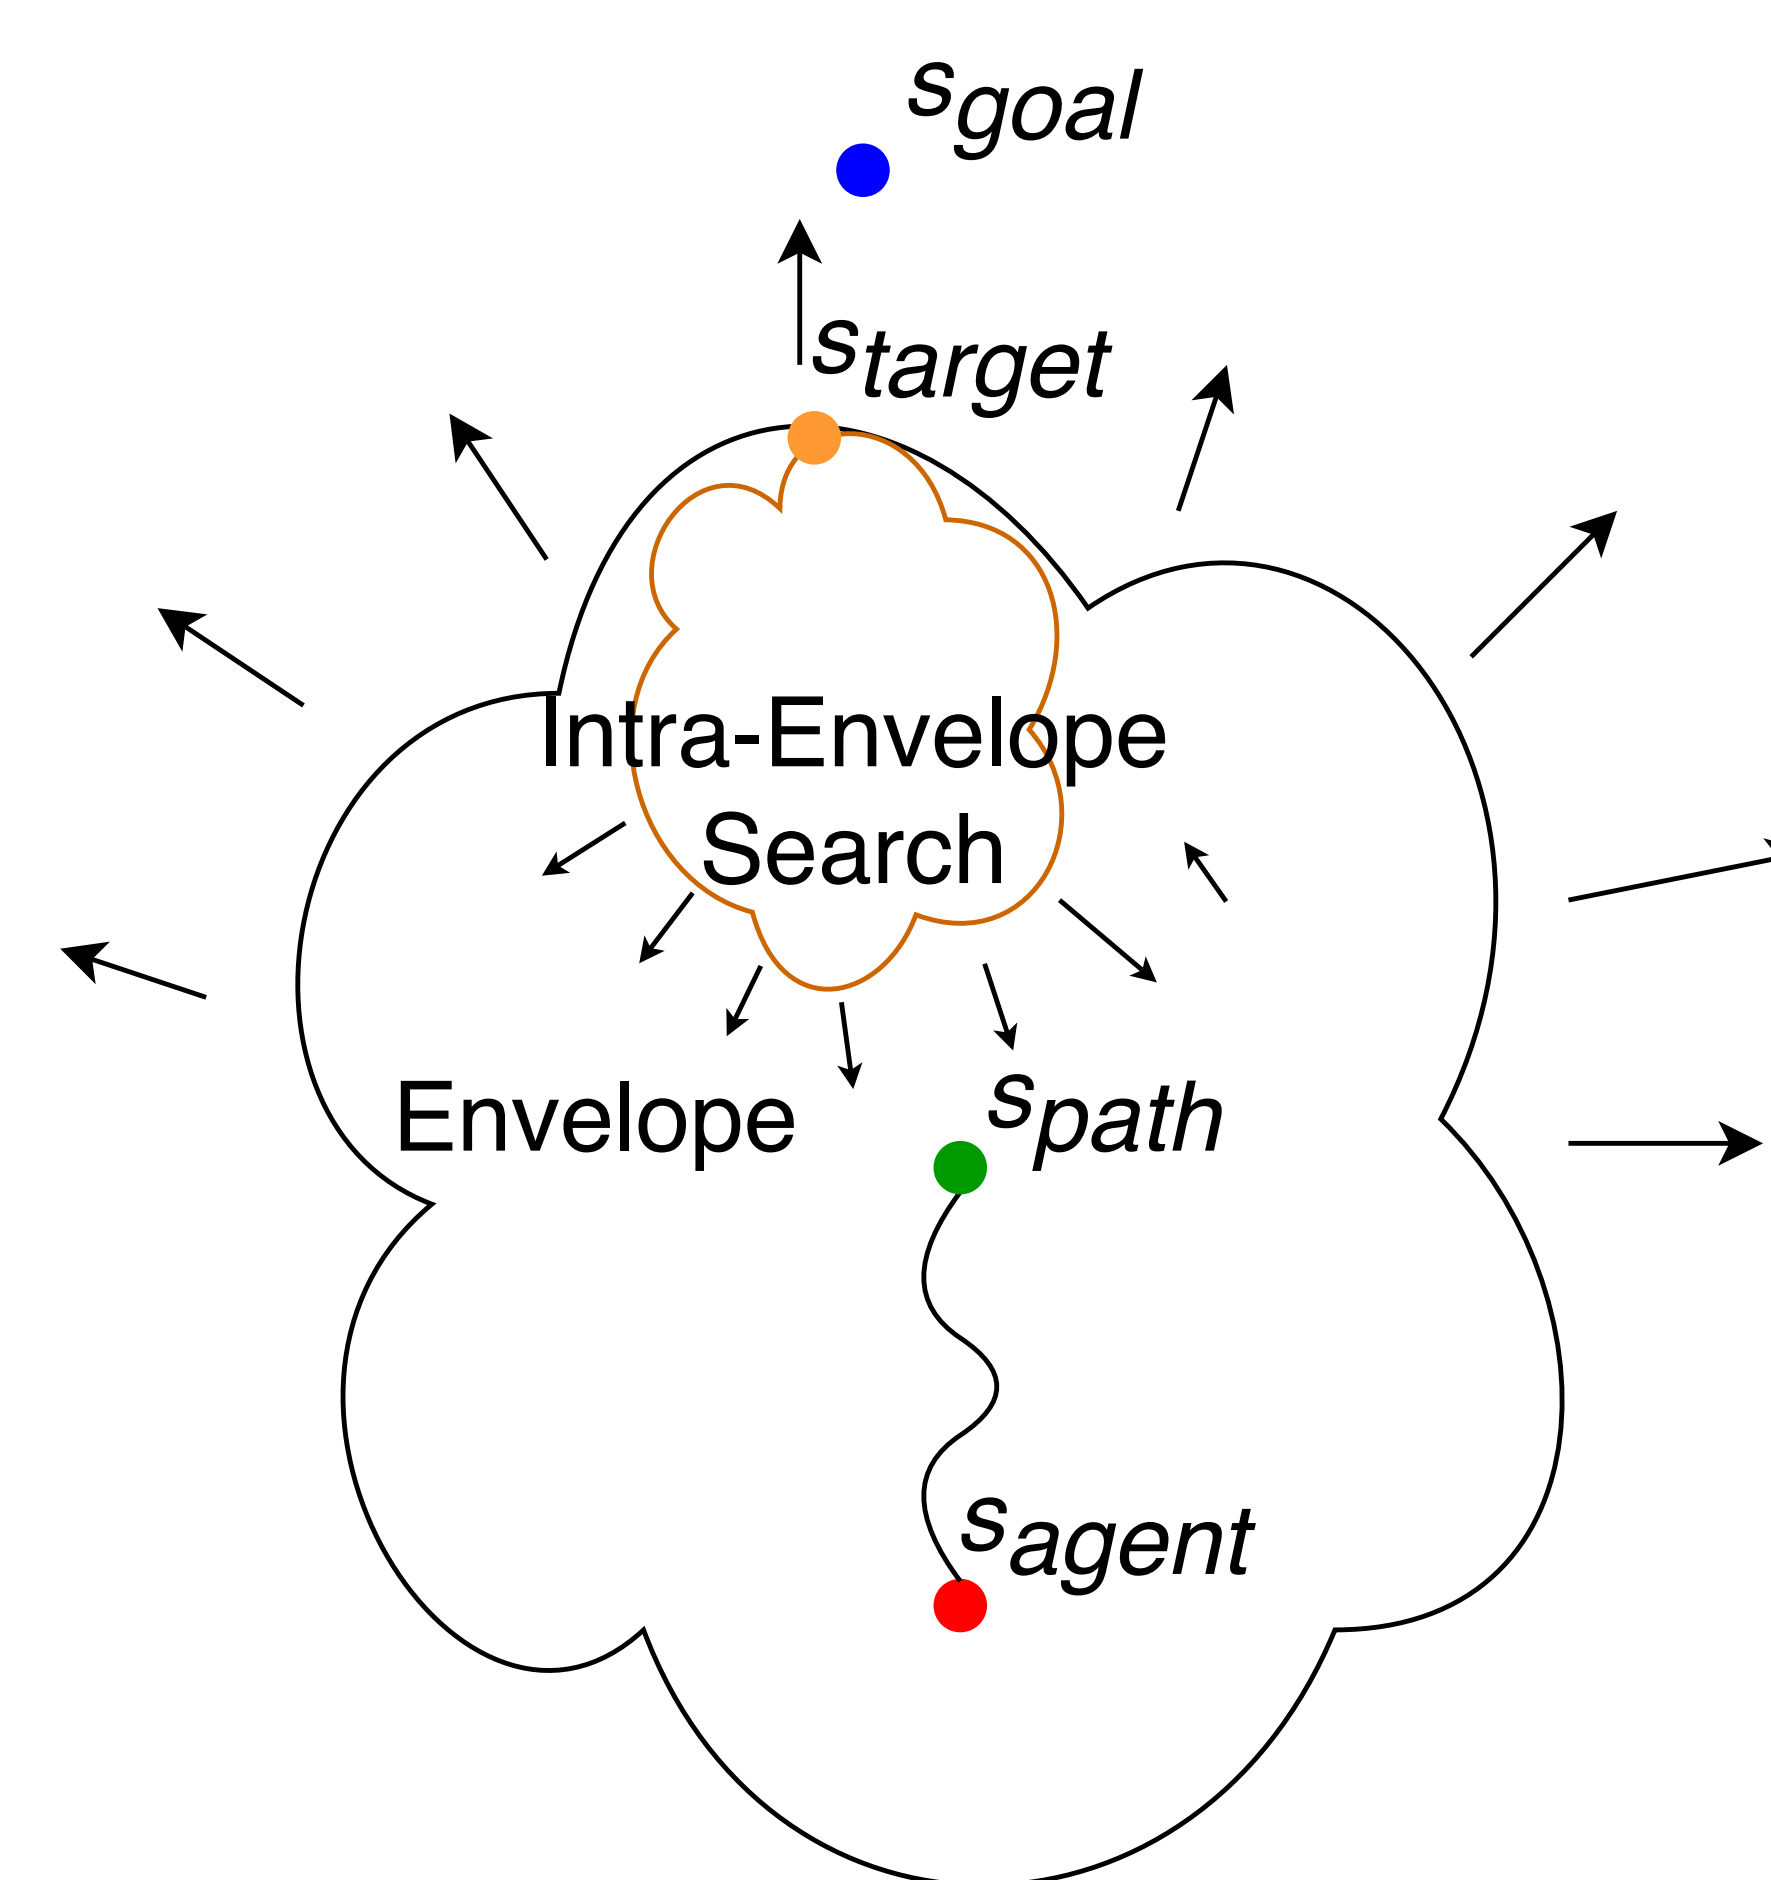




Illustration of I-ES. A search is performed within the envelope from some target to the agent’s cached path.



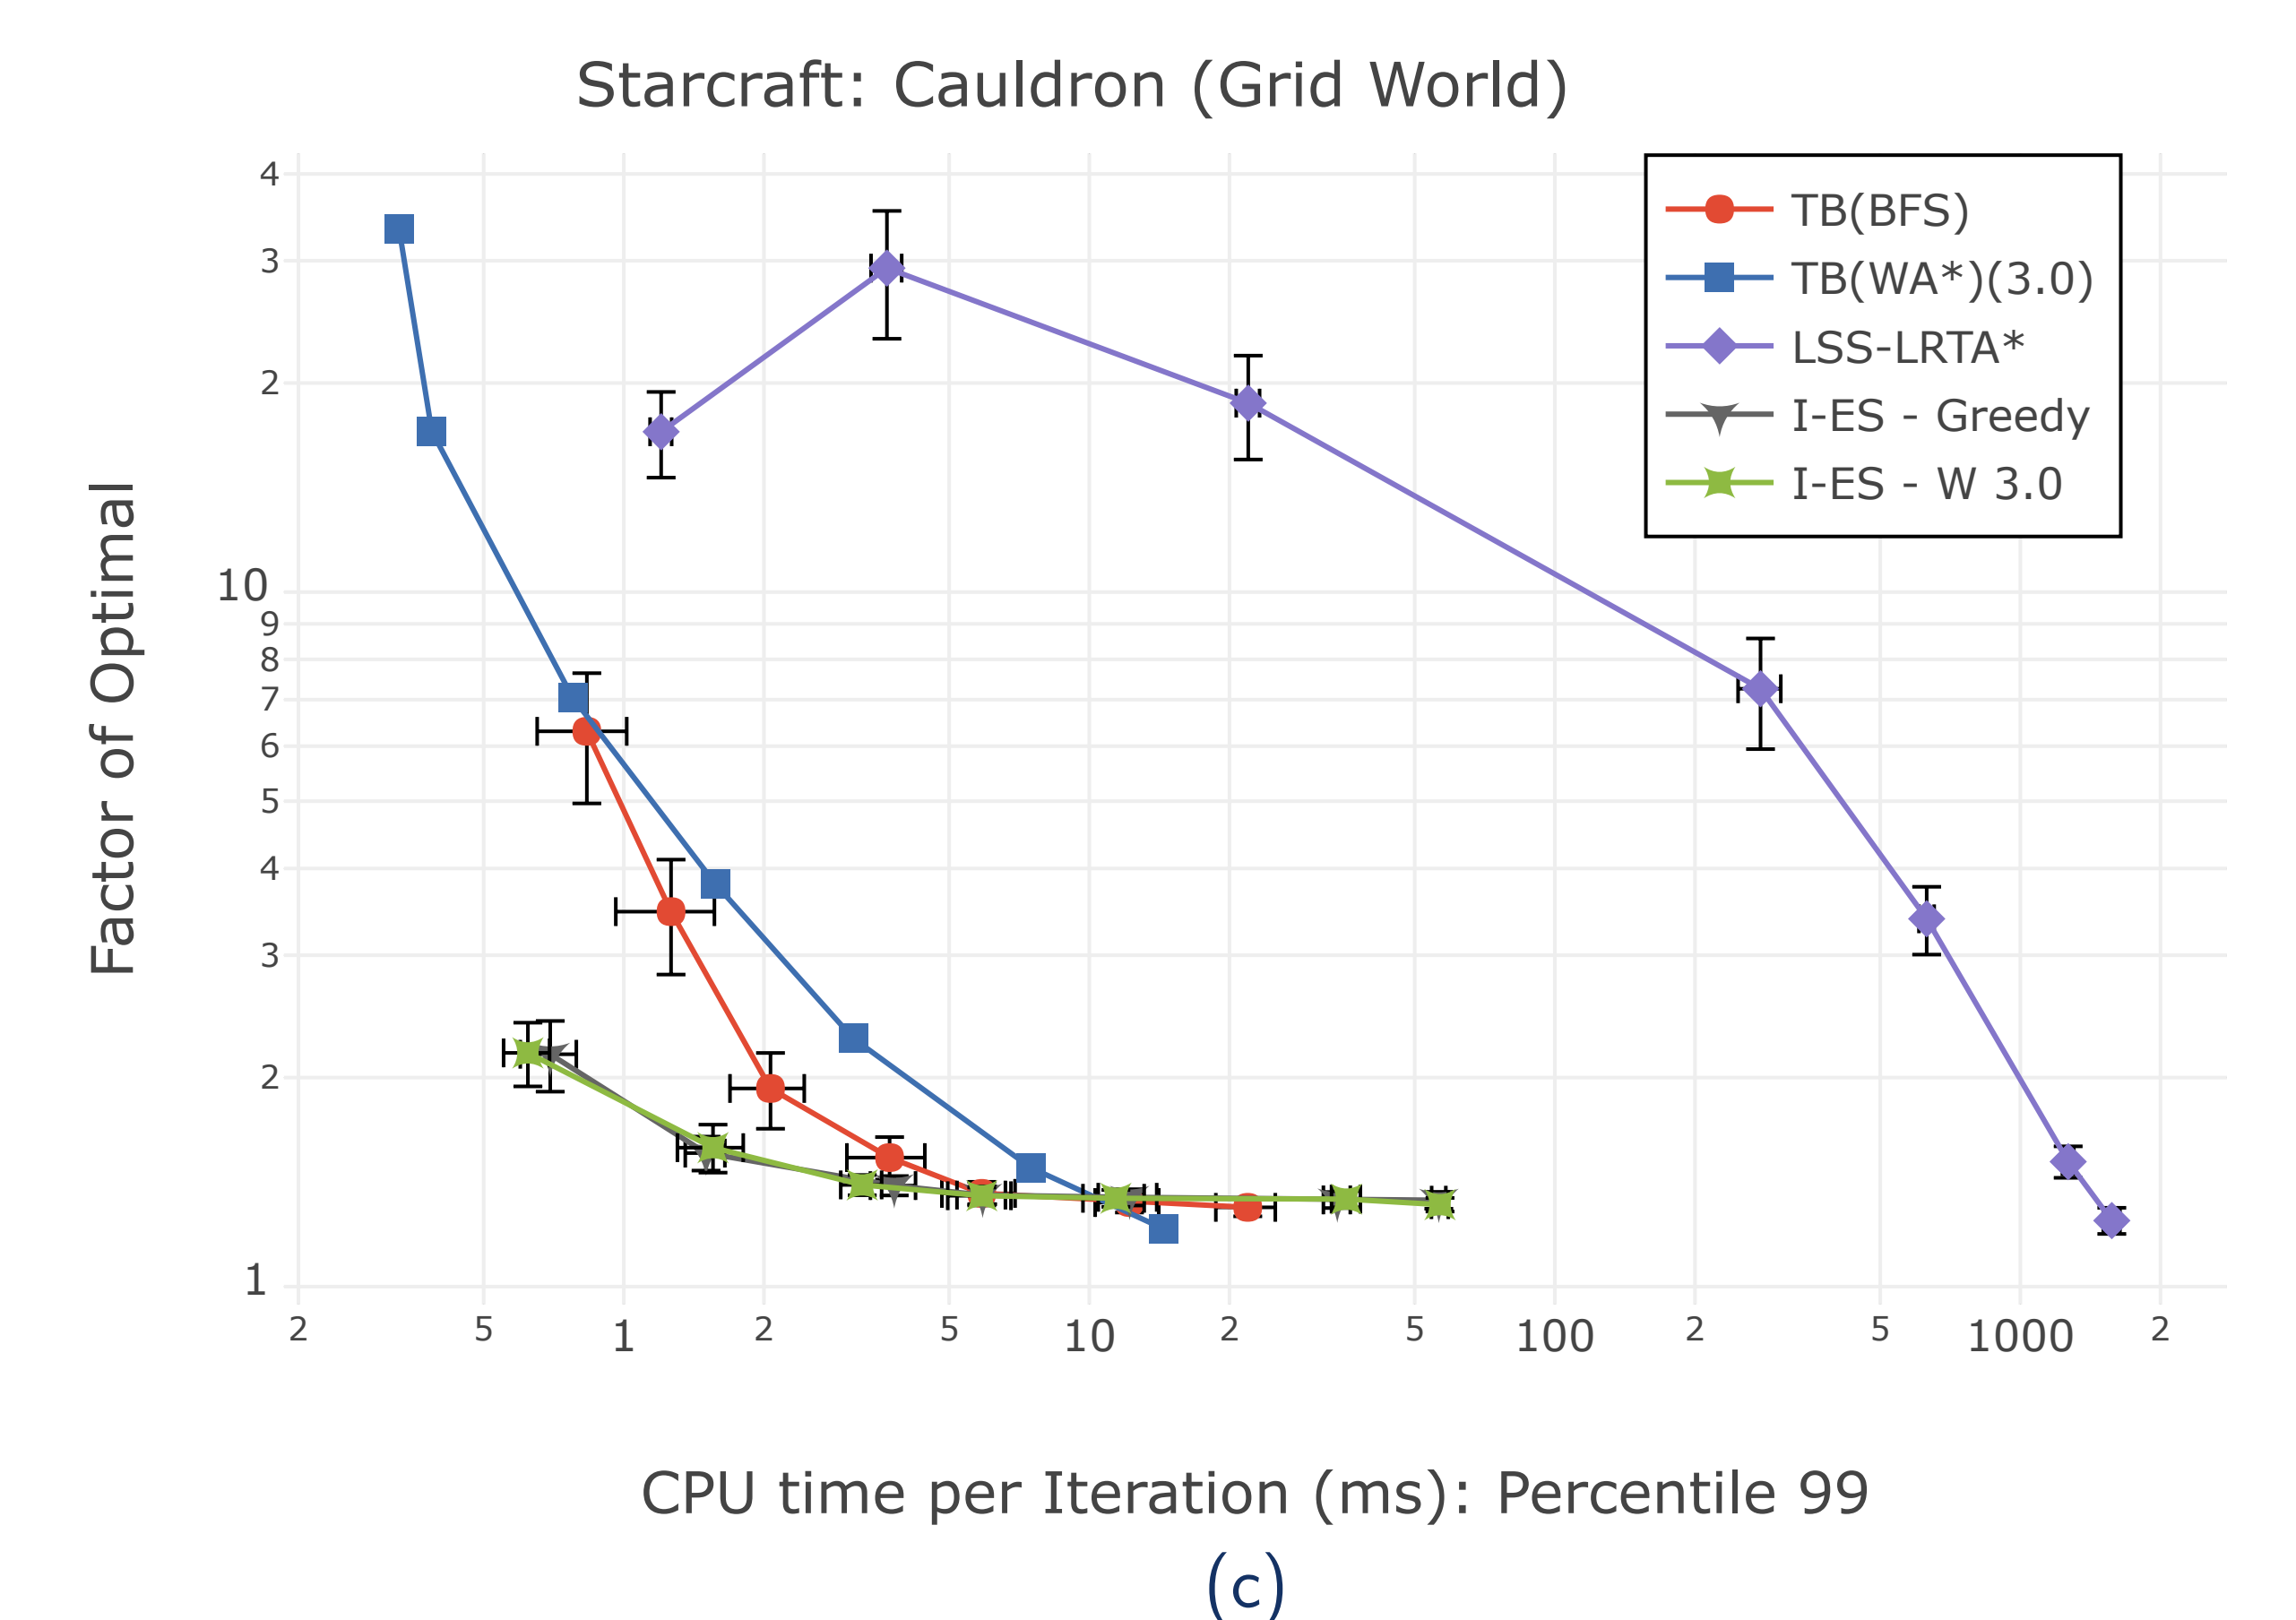

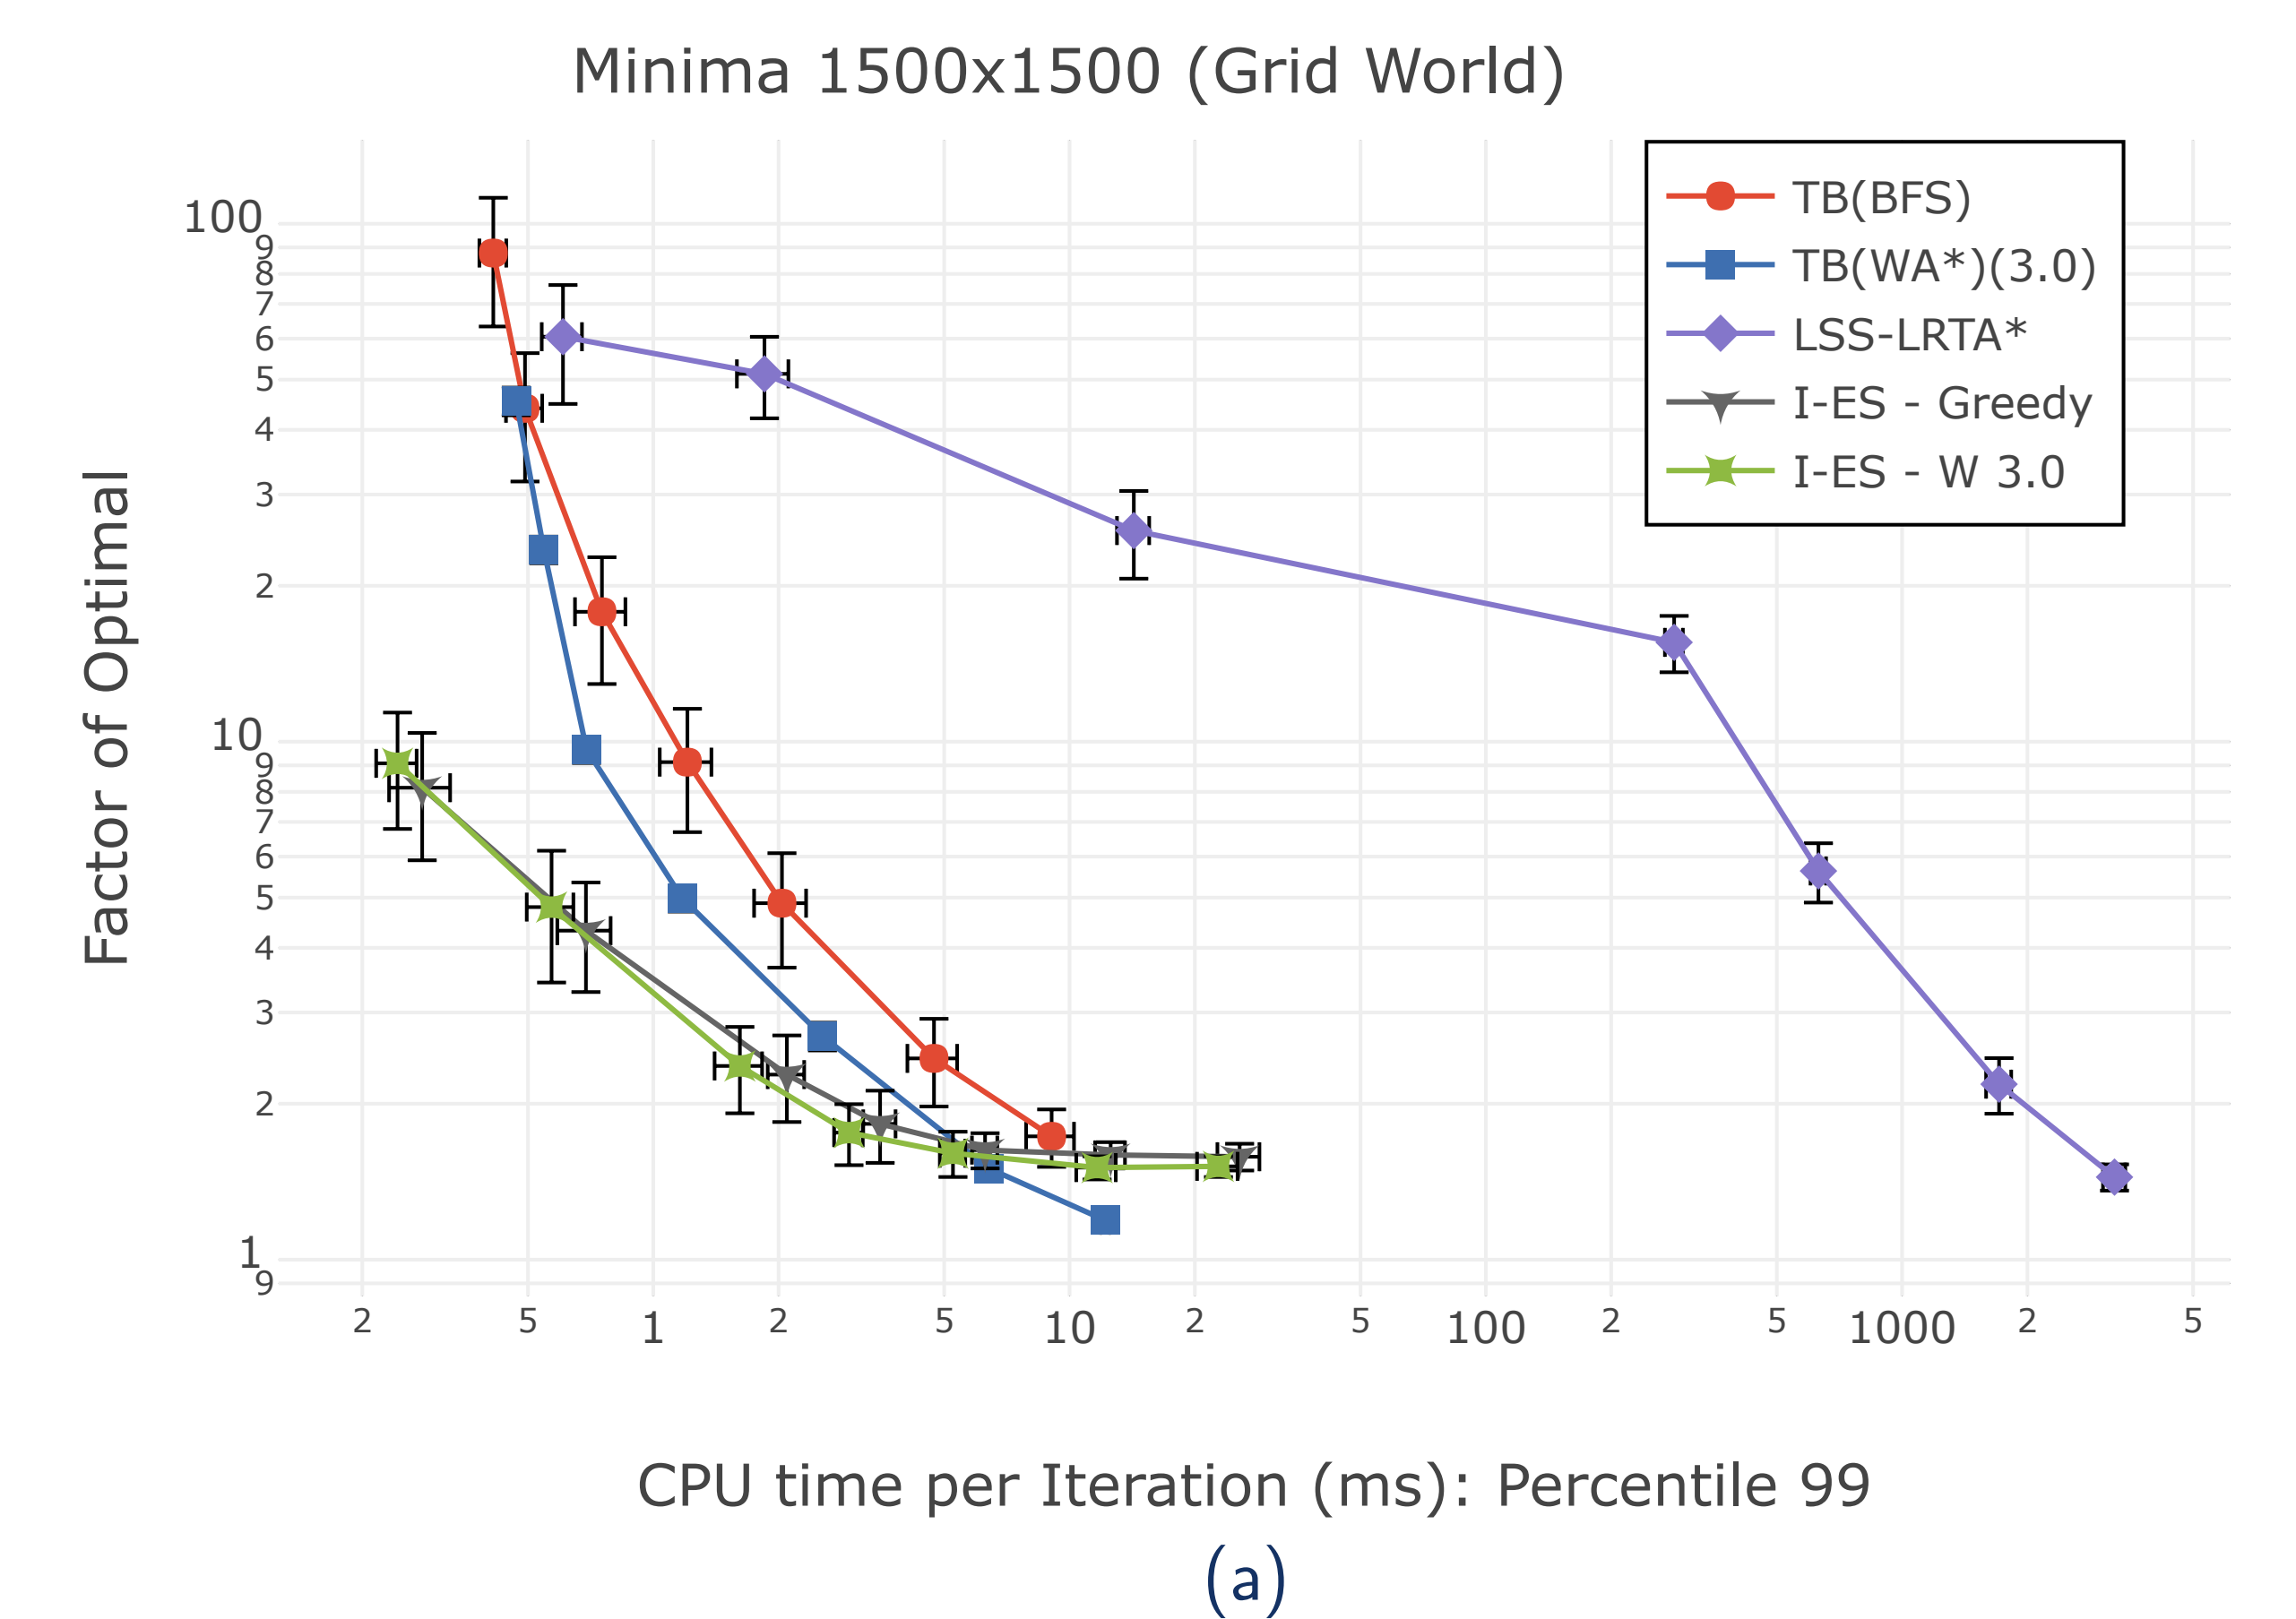
I-ES and $TB(WA^*)$ explore a sub-optimal branch (red) by following the heuristic. Restriction to A* tree forces $TB(WA^*)$ to backtrack (green) while I-ES continues along the shortcut path (blue) to the goal.



I-ES and $TB(BFS)$ must expand every state in the room after entering it (red path). When the exit node is discovered, $TB(BFS)$ backtracks (green) while I-ES follows the shortcut (blue) to the exit.

Empirical Evaluation

- Time bounds specified as limits on number of expanded nodes, reported as CPU time per iteration
- y-axis: mean agent trajectory cost as factor of A* offline optimal solution; x-axis: 99th percentile of CPU time per iteration



Conclusions

- I-ES is State-of-the-Art in grids with tight time bounds
- Overhead is higher than TBA*, but worthwhile for dense graphs
- Envelope-based algorithms deserve more attention: real-time search without scrubbing!