Planning with Experience Graphs

Mike Phillips
Carnegie Mellon University

Collaborators
Benjamin Cohen, Andrew Dornbush, Victor Hwang,
Sachin Chitta, Maxim Likhachev

Motivation

Many tasks are repetitive. They may have different starts and goals, but have the same general motion.

Examples:

- loading a dishwasher
- opening doors
- moving objects around a warehouse

- **Robots should be able to re-use prior experience to accelerate planning**
- **Especially useful for high-dimensional planning problems such as mobile manipulation**

Background

- Find a collision free, good quality, path from the start state to the goal state

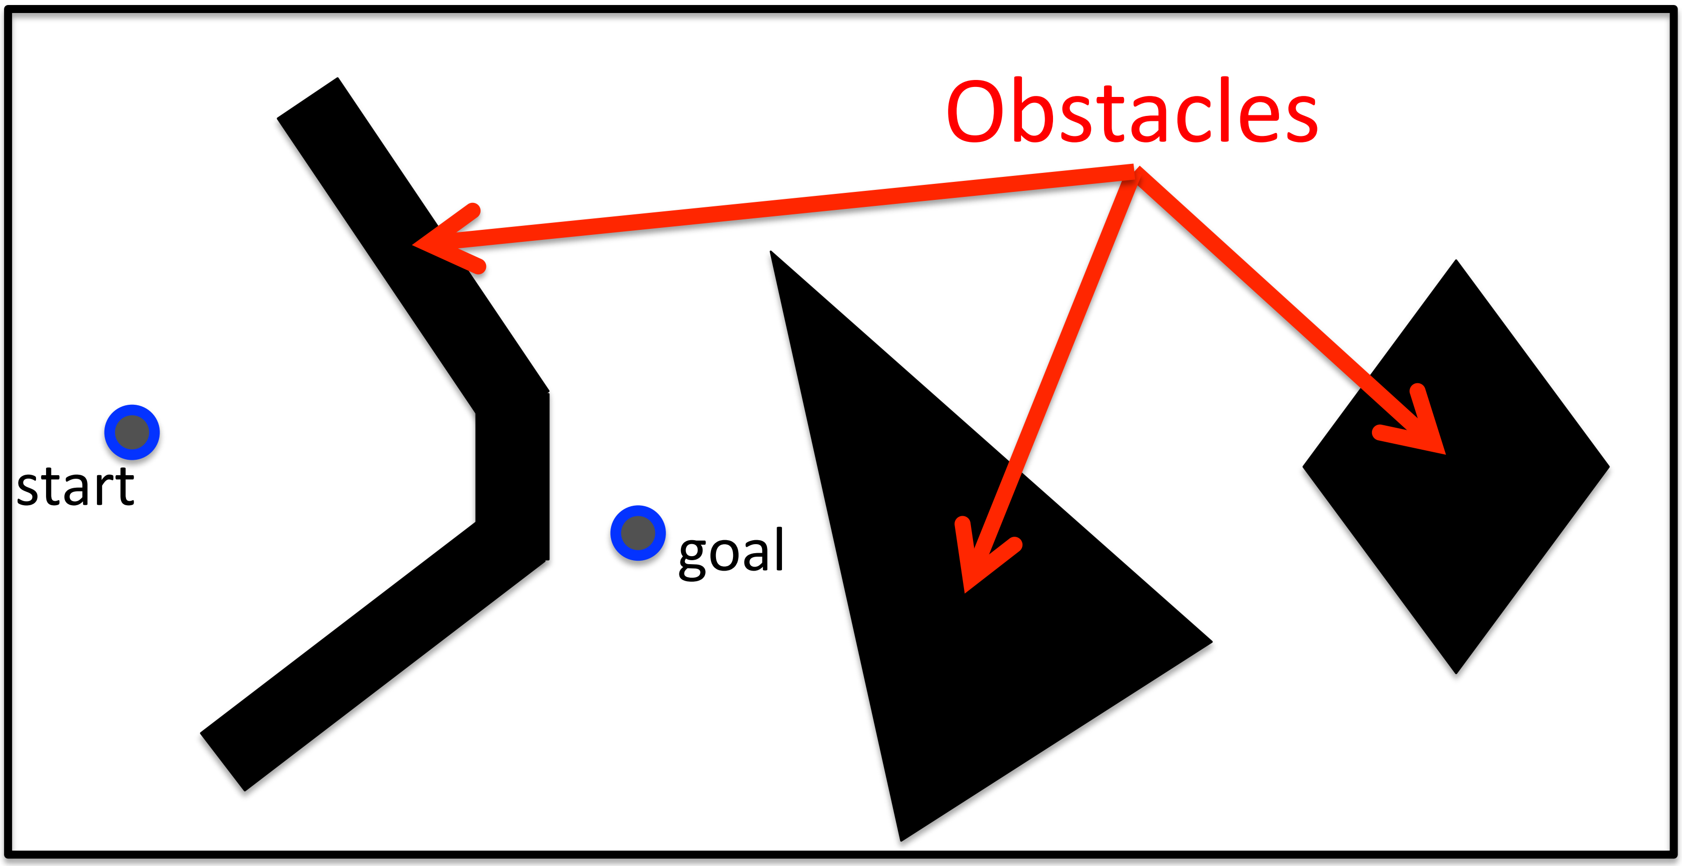

Background

Graph representing free space

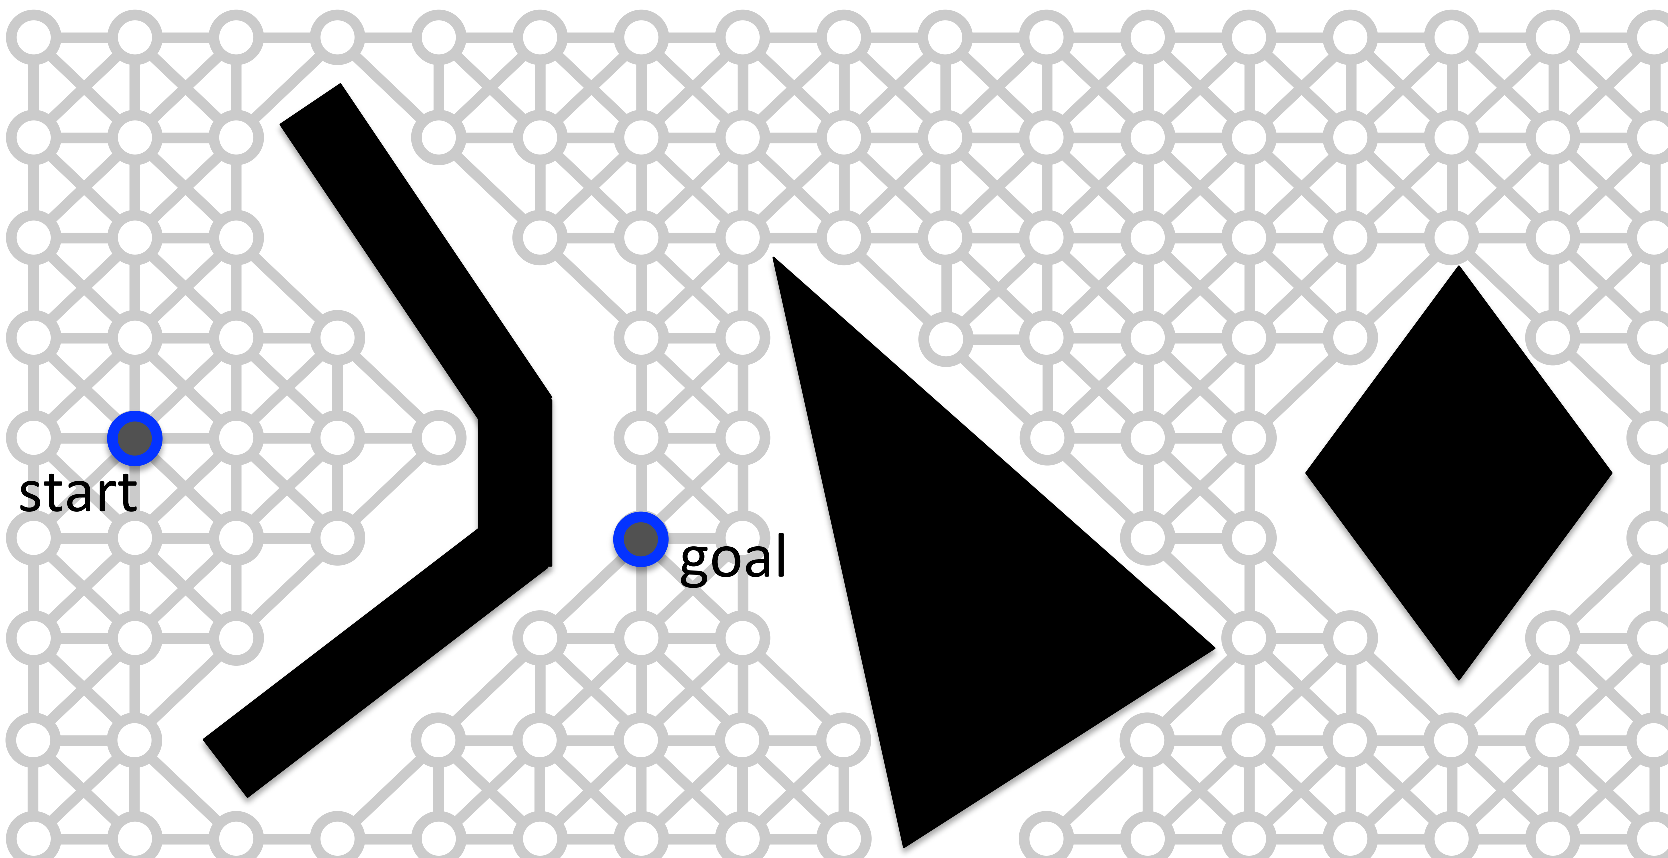

Background

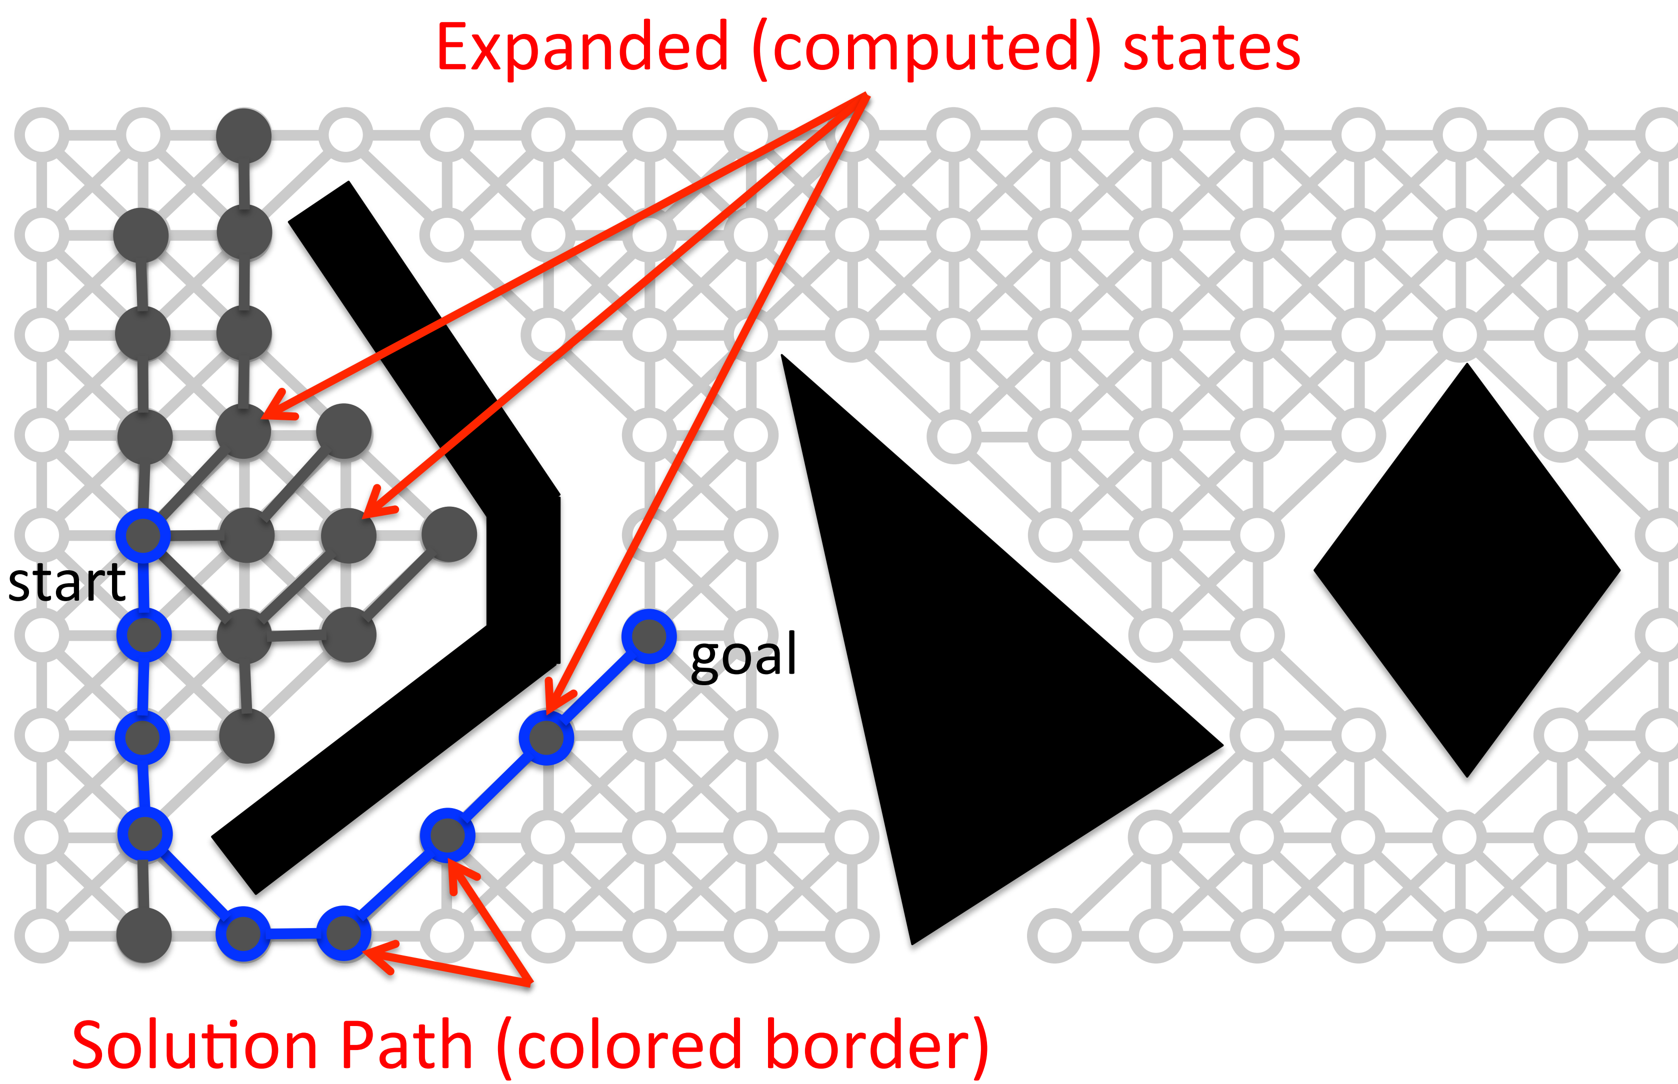
Expanded (computed) states
start
goal
Solution Path (colored border)

Background

- A similar scenario
- This repeats a lot of computation!

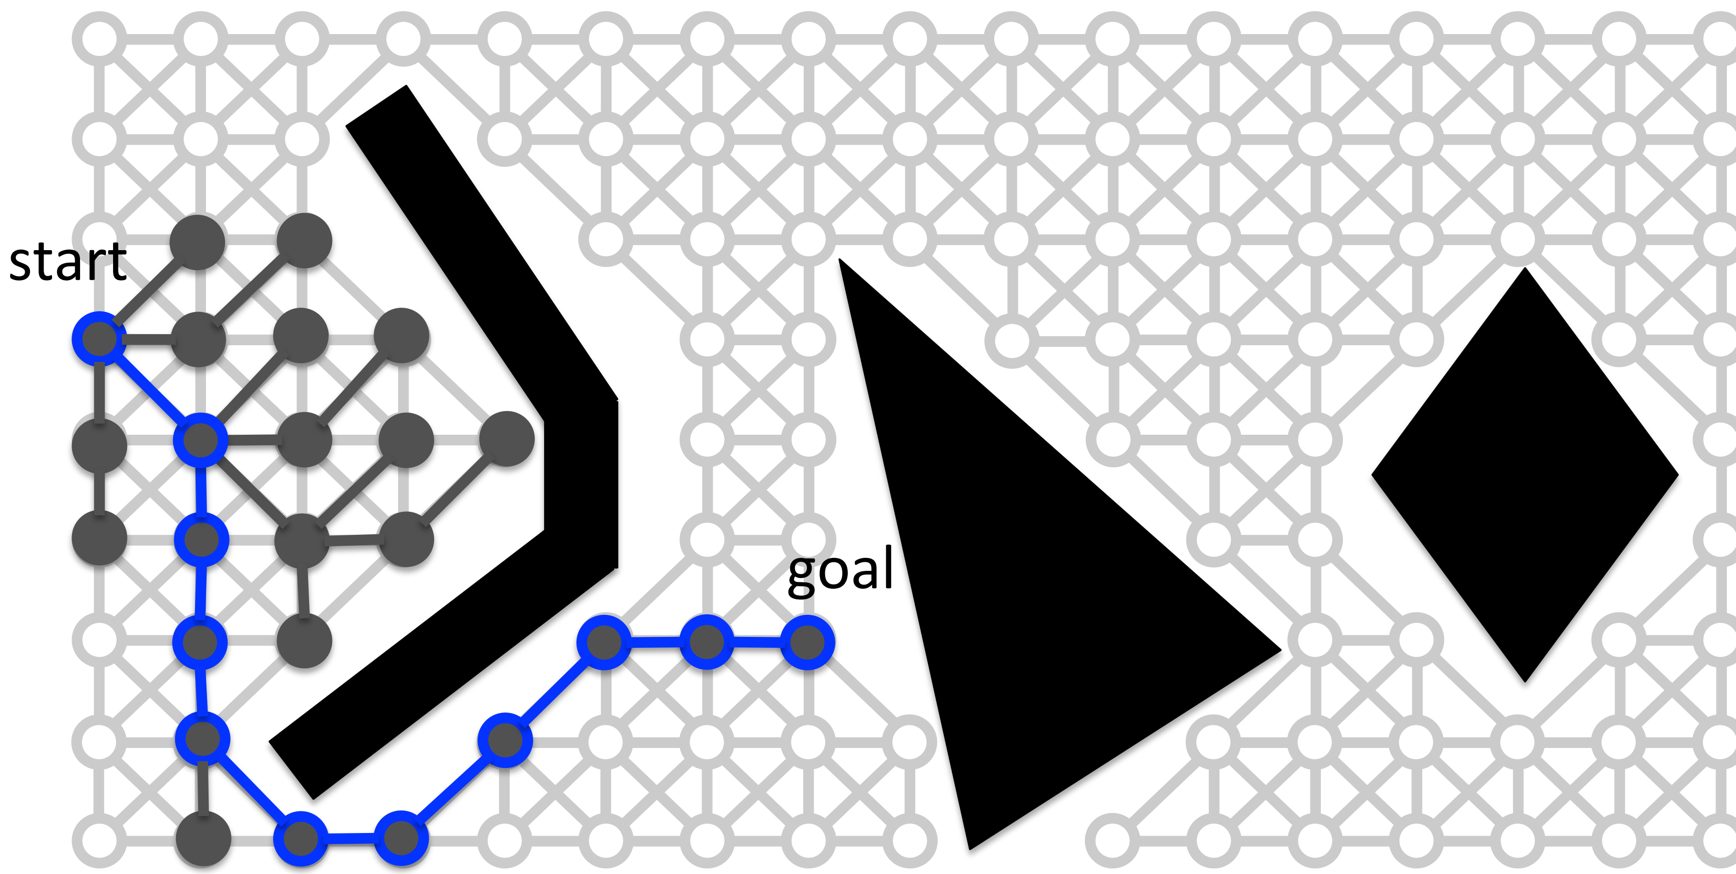

Experience Graphs (E-Graphs)

- Collection of previously computed paths or demonstrations
- A significantly smaller sub-graph of the original graph

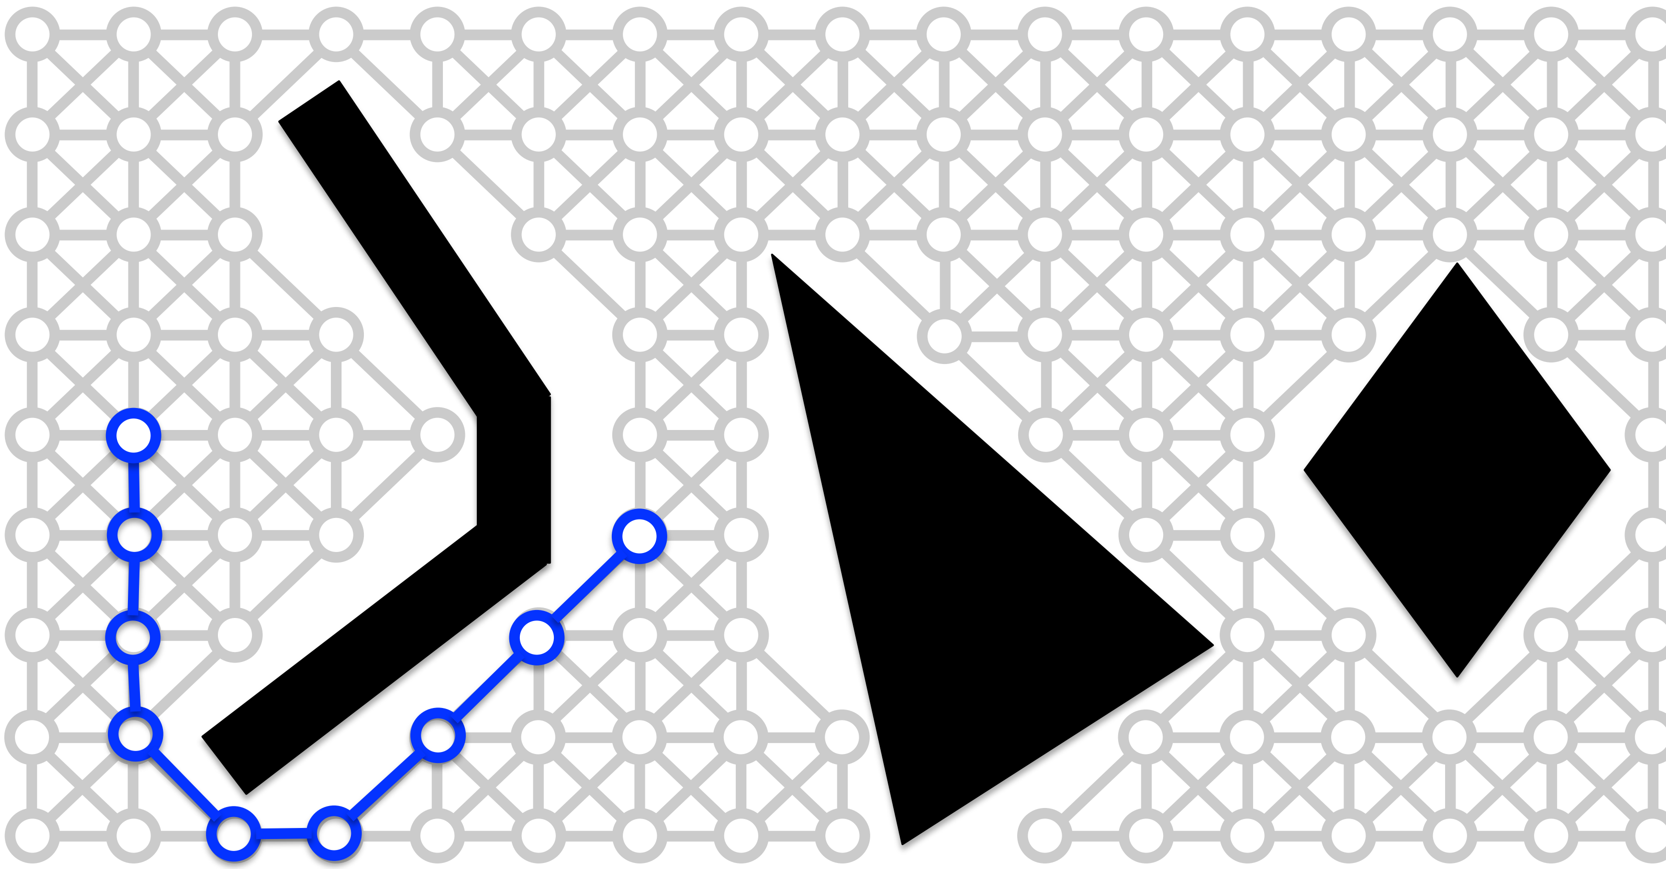

Experience Graphs (E-Graphs)

- For repetitive tasks, planning with E-Graphs is much faster

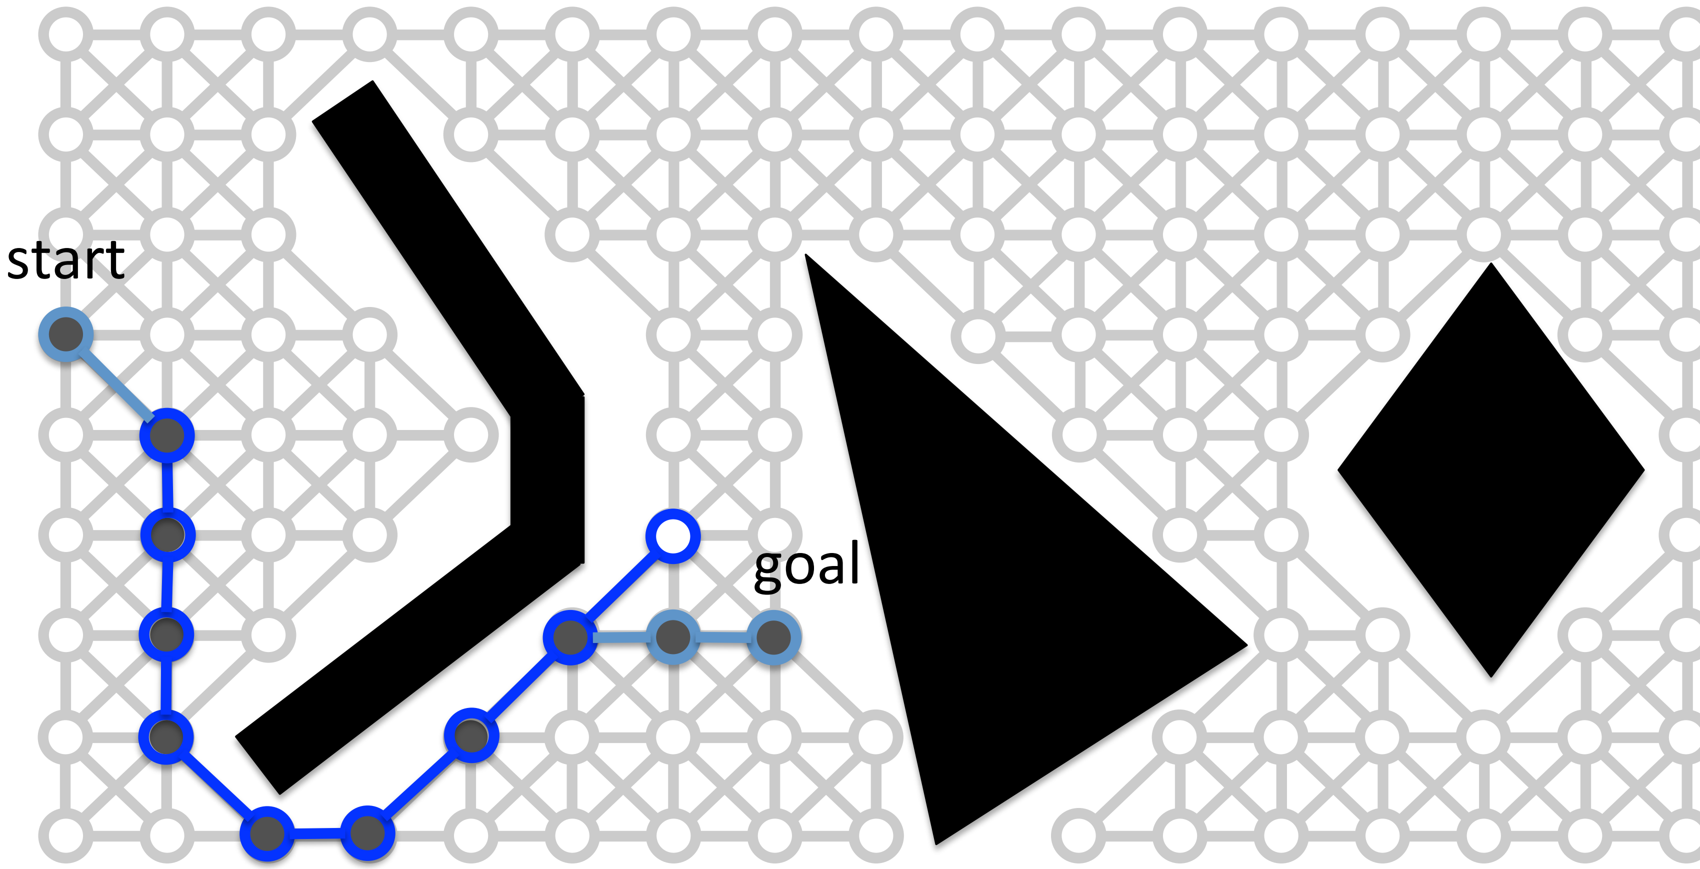

Experience Graphs (E-Graphs)

- For repetitive tasks, planning with E-Graphs is much faster

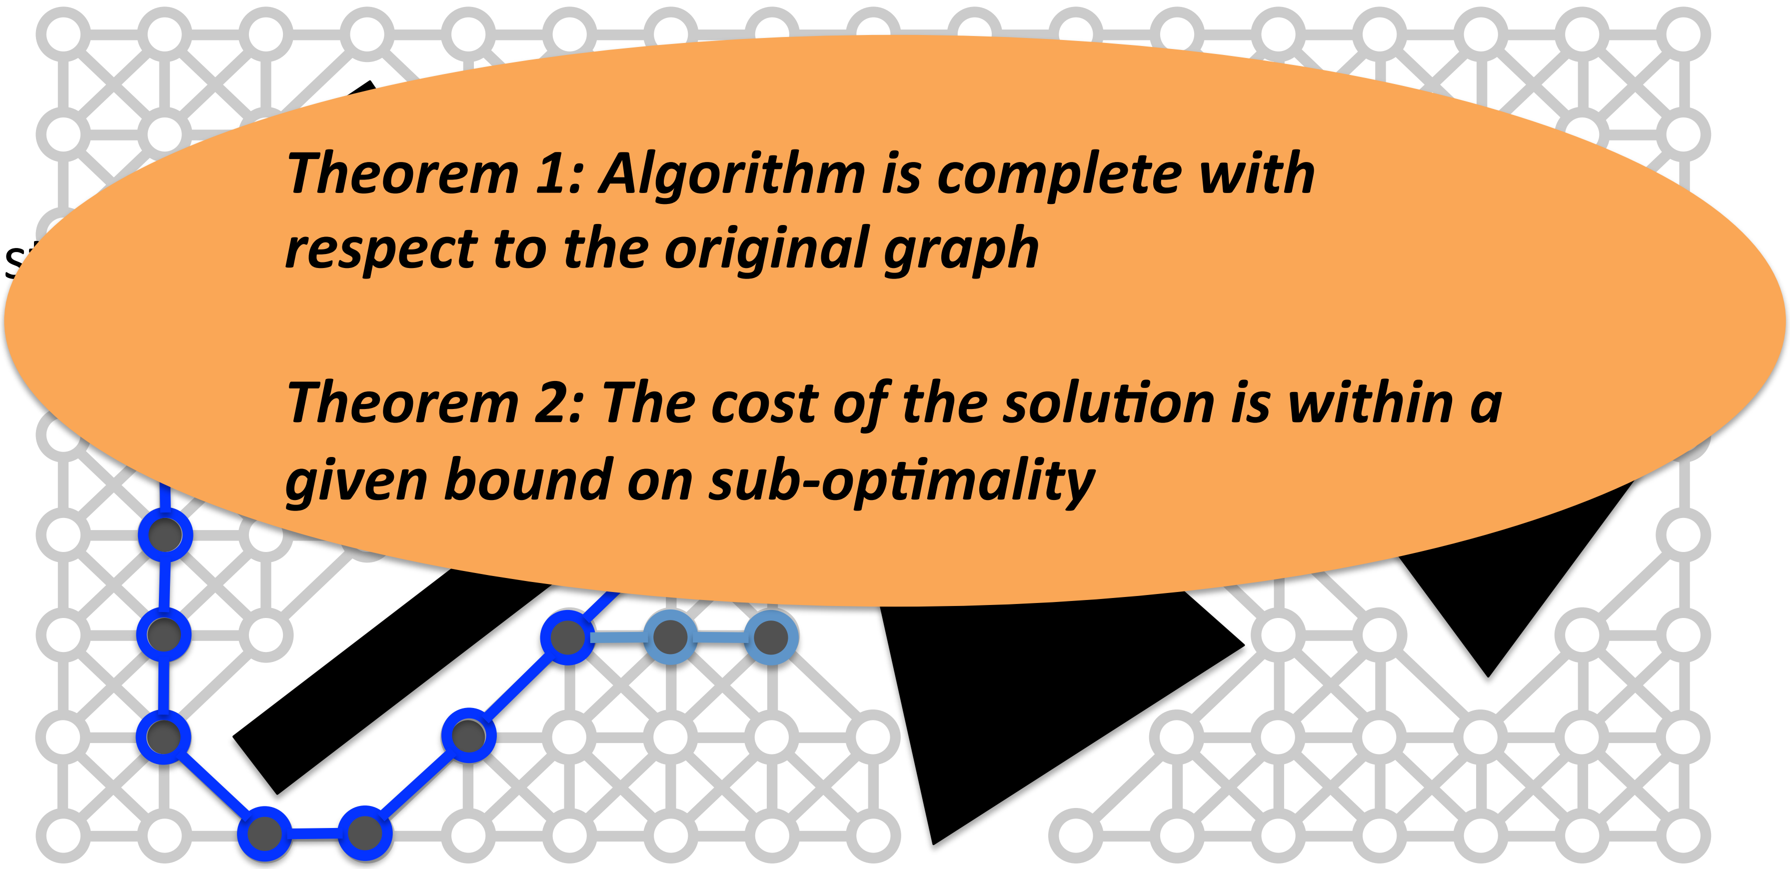

Theorem 1: Algorithm is complete with respect to the original graph

Theorem 2: The cost of the solution is within a given bound on sub-optimality

Experience Graphs (E-Graphs)

- E-Graph after three runs…

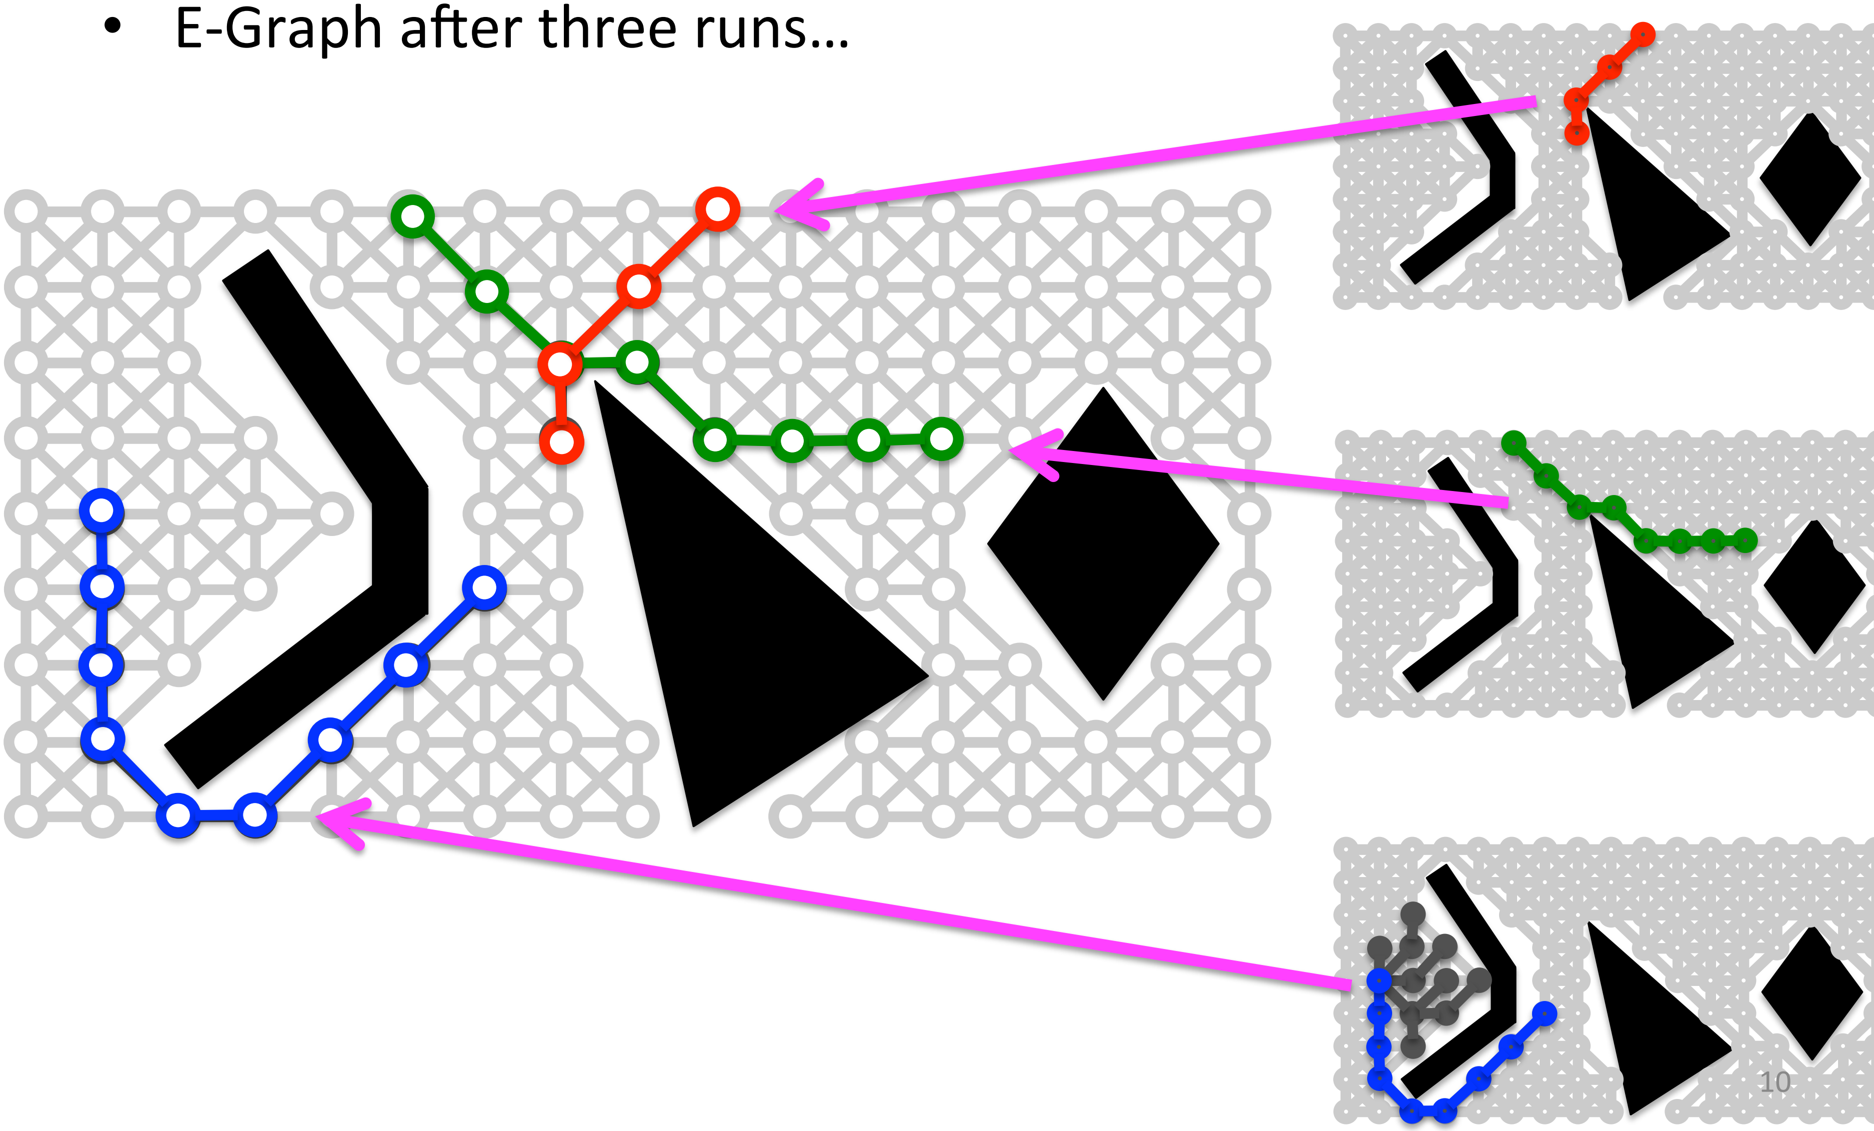

Experience Graphs (E-Graphs)

- E-Graph after three runs…

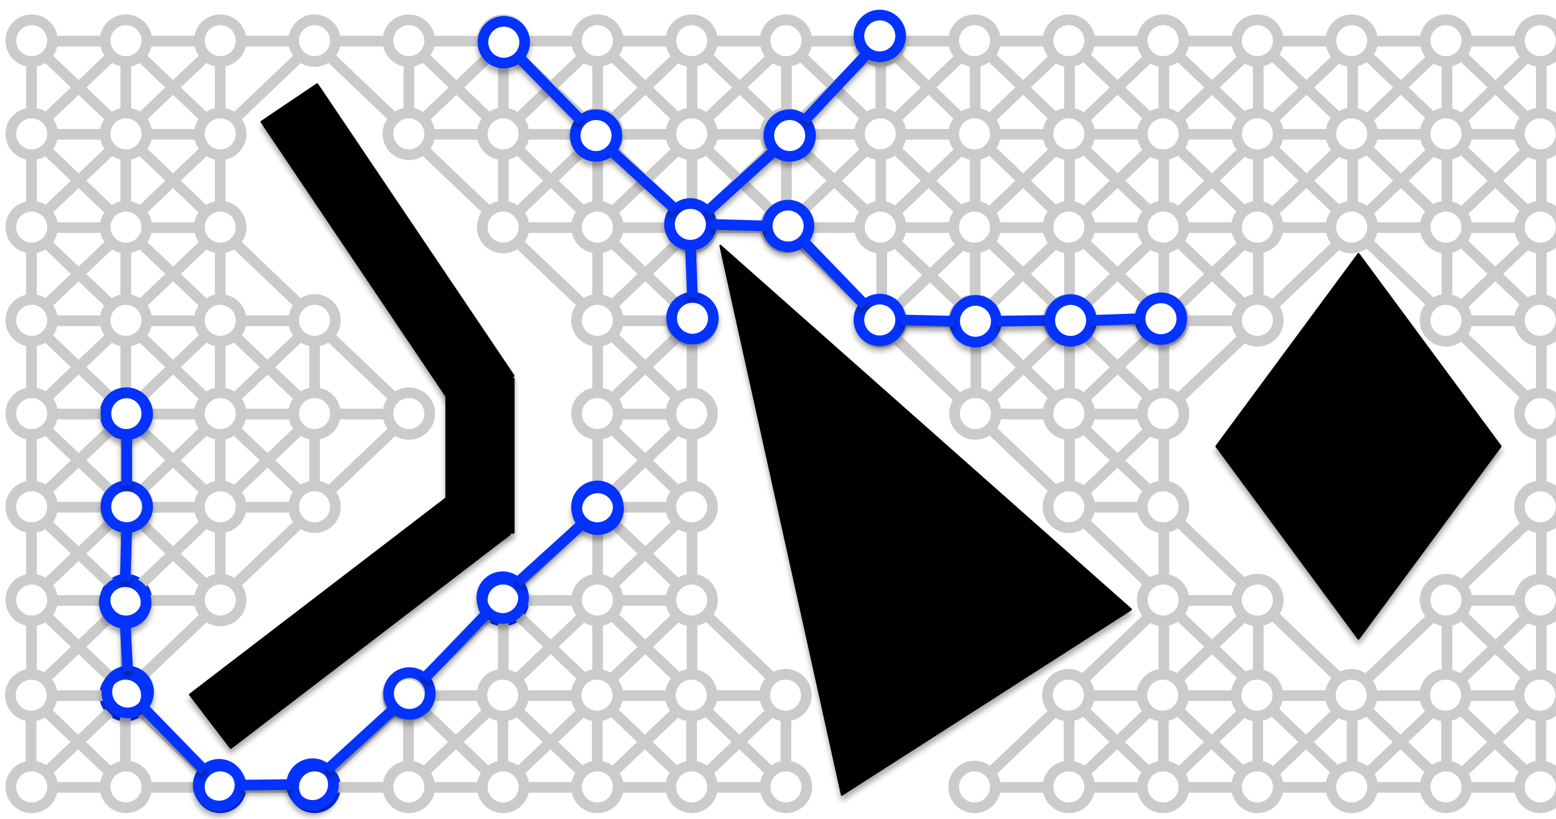

Experience Graphs (E-Graphs)

- Using E-graph
  - Very few states expanded
  - **Completeness & bounds on sub-optimality w.r.t. original graph**

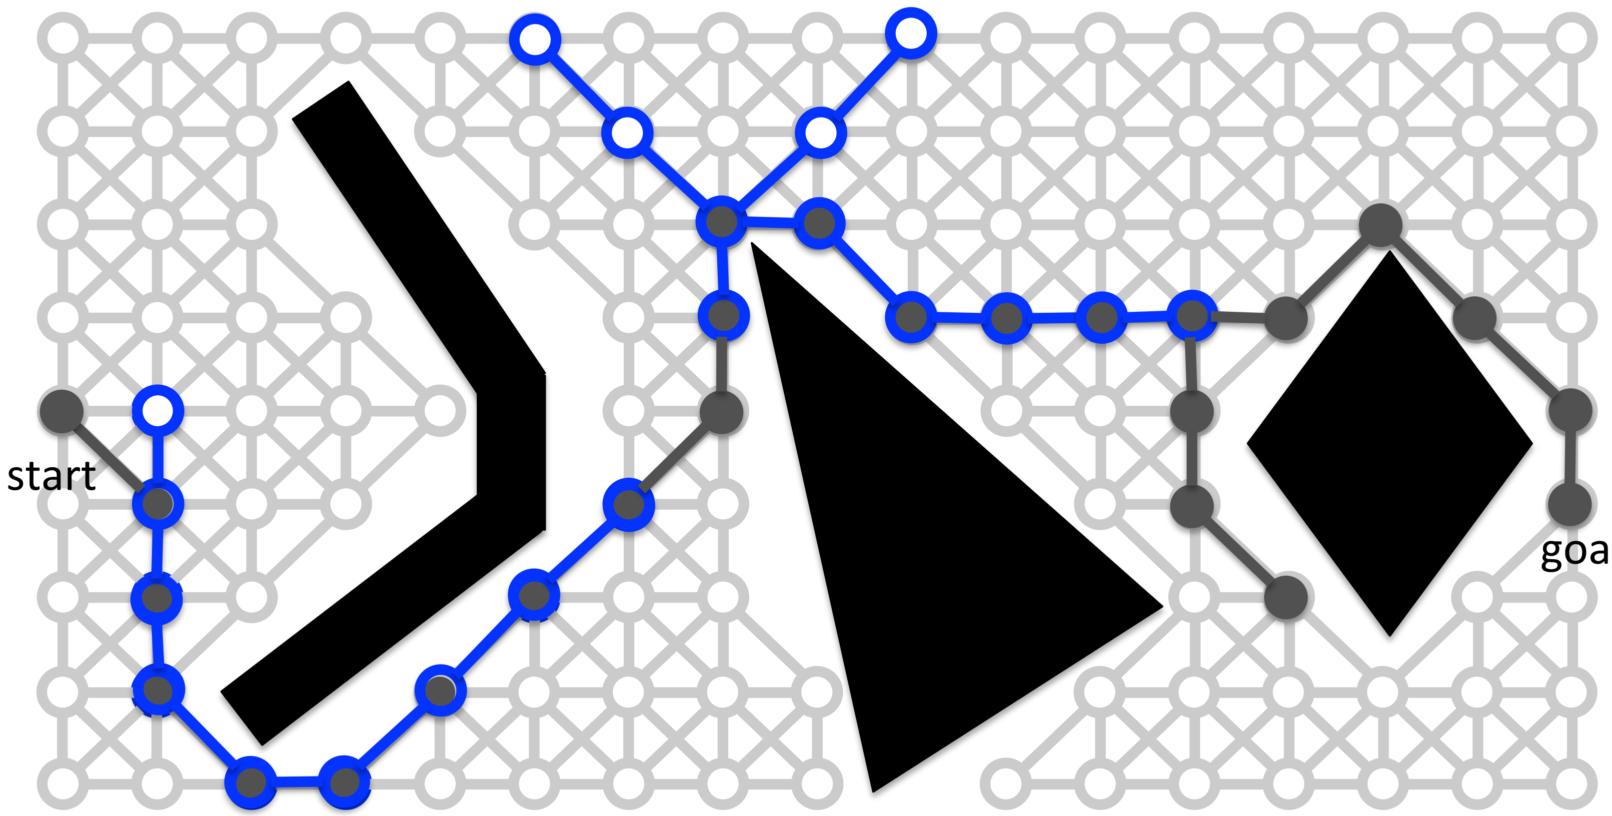

Experience Graphs (E-Graphs)

- Reuse E-Graph by:
  - Introducing a new heuristic function
  - Heuristic guides the search toward expanding states on the E-Graph

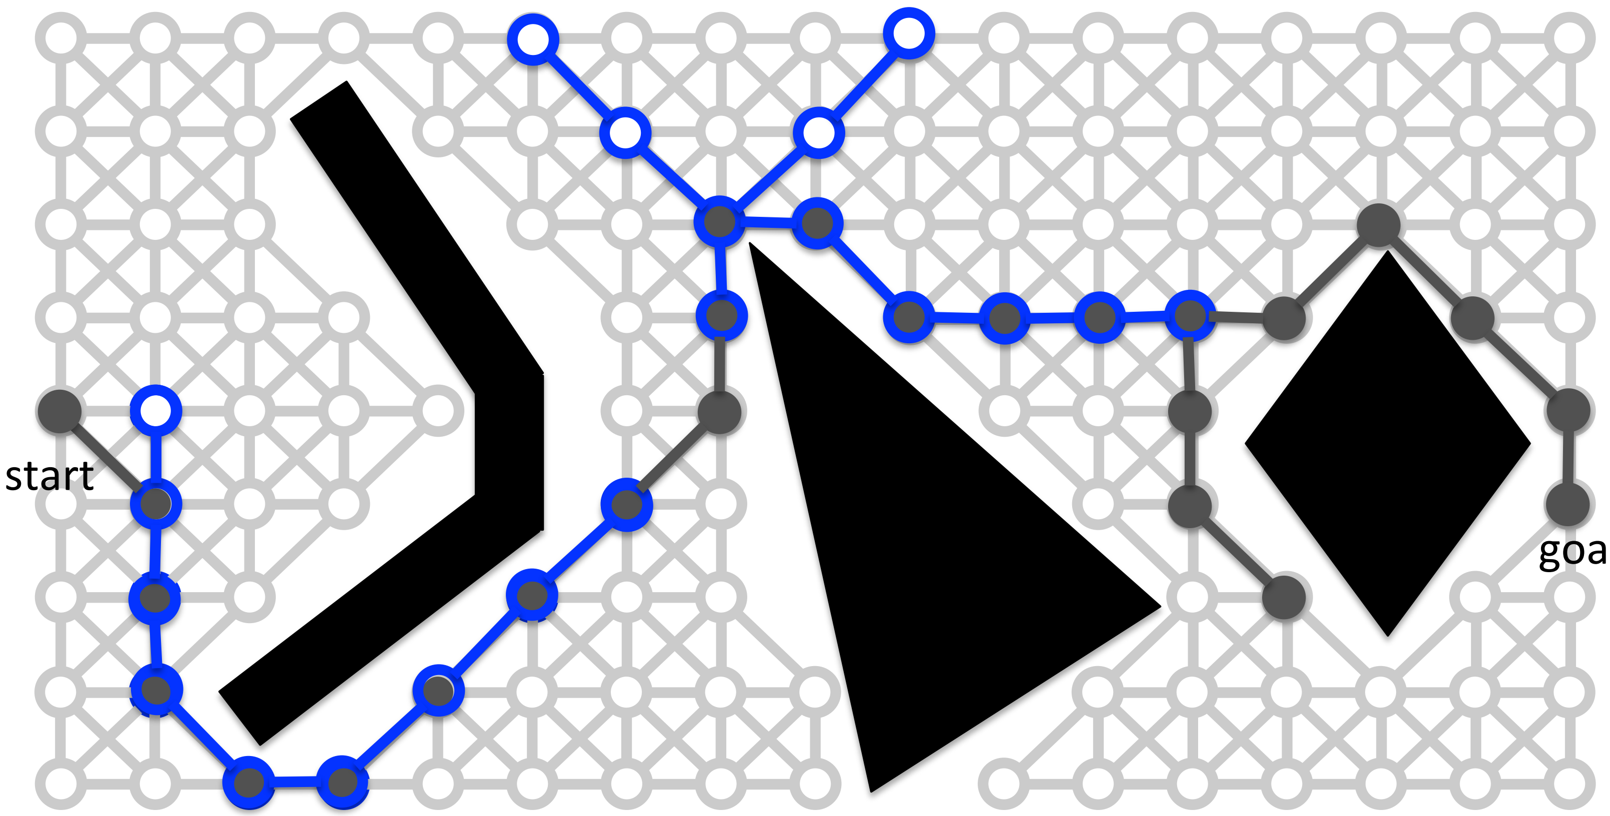

Heuristic

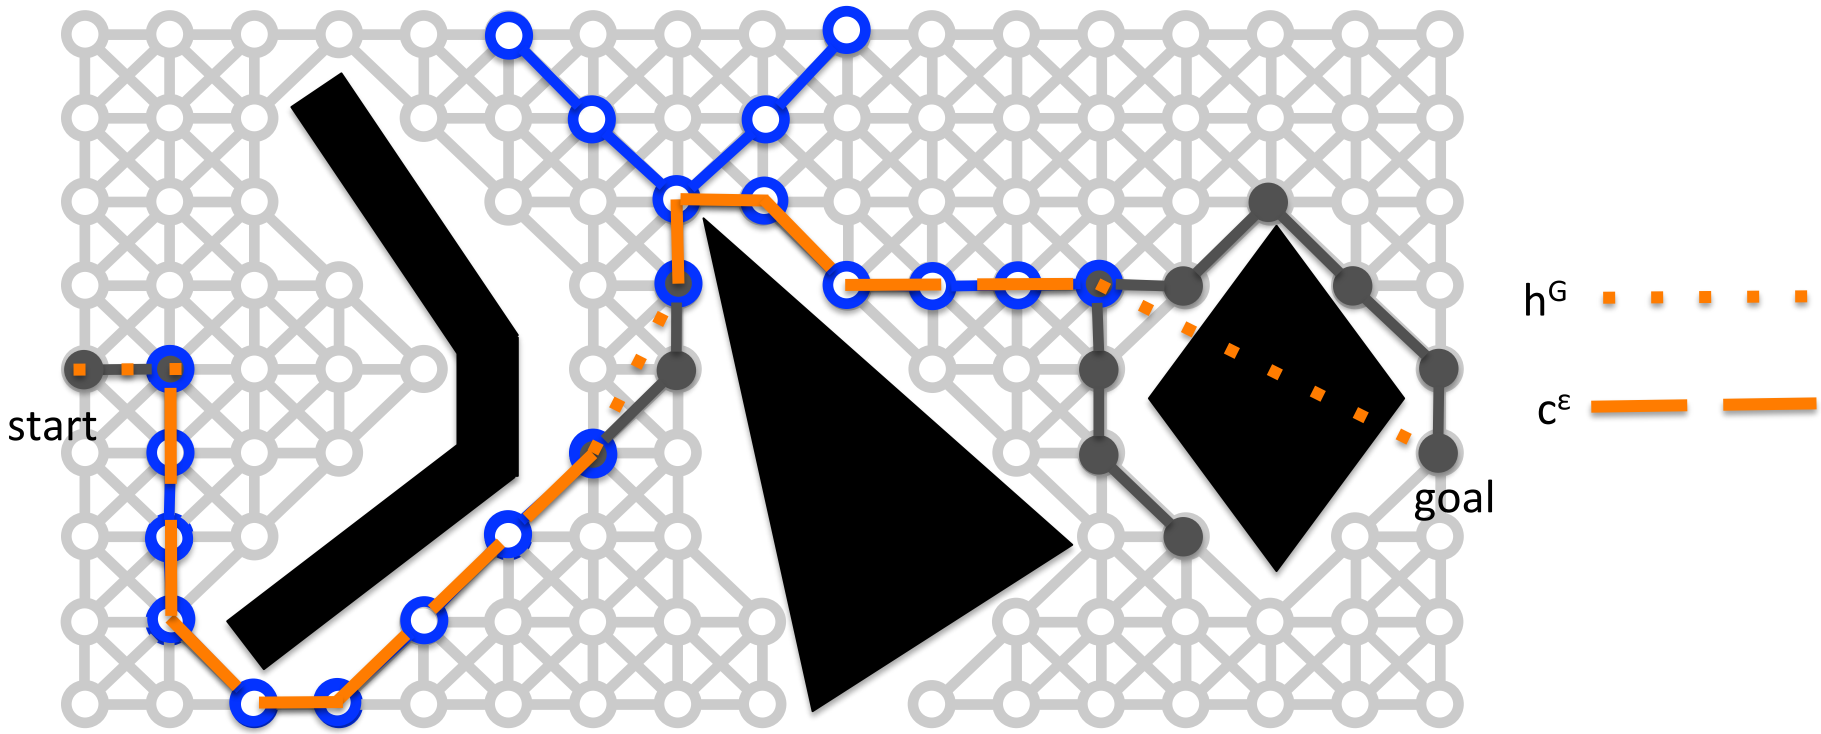

Heuristic computation finds a min cost path using two kinds of "edges"

Travelling off the E-Graph uses an inflated original heuristic

Travelling on E-Graph uses actual costs

$$h^{\mathcal{E}}(s_0) = \min_{\pi} \sum_{i=0}^{N-1} \min\{\varepsilon^{\mathcal{E}} h^G(s_i, s_{i+1}), c^{\mathcal{E}}(s_i, s_{i+1})\}$$

Heuristic

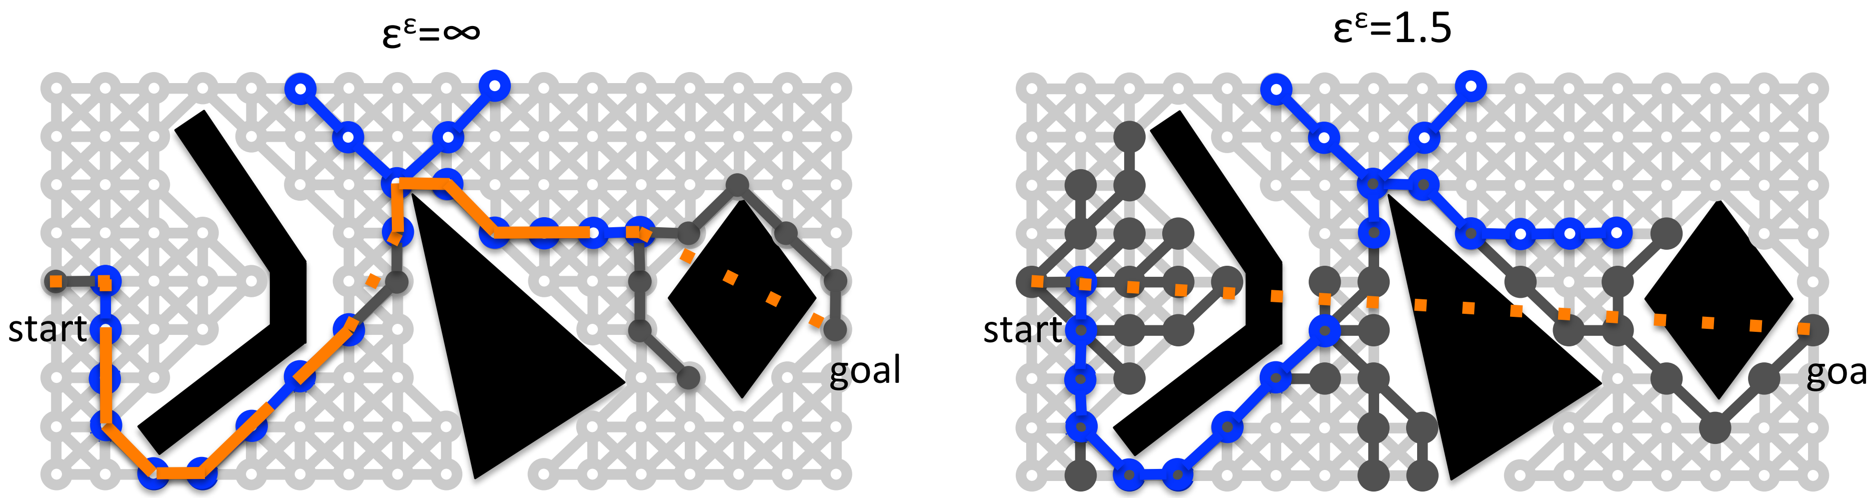

"E-Graphs: Bootstrapping Planning with Experience Graphs"
Mike Phillips, Benjamin Cohen, Sachin Chitta, Maxim Likhachev
RSS 2012

Heuristic

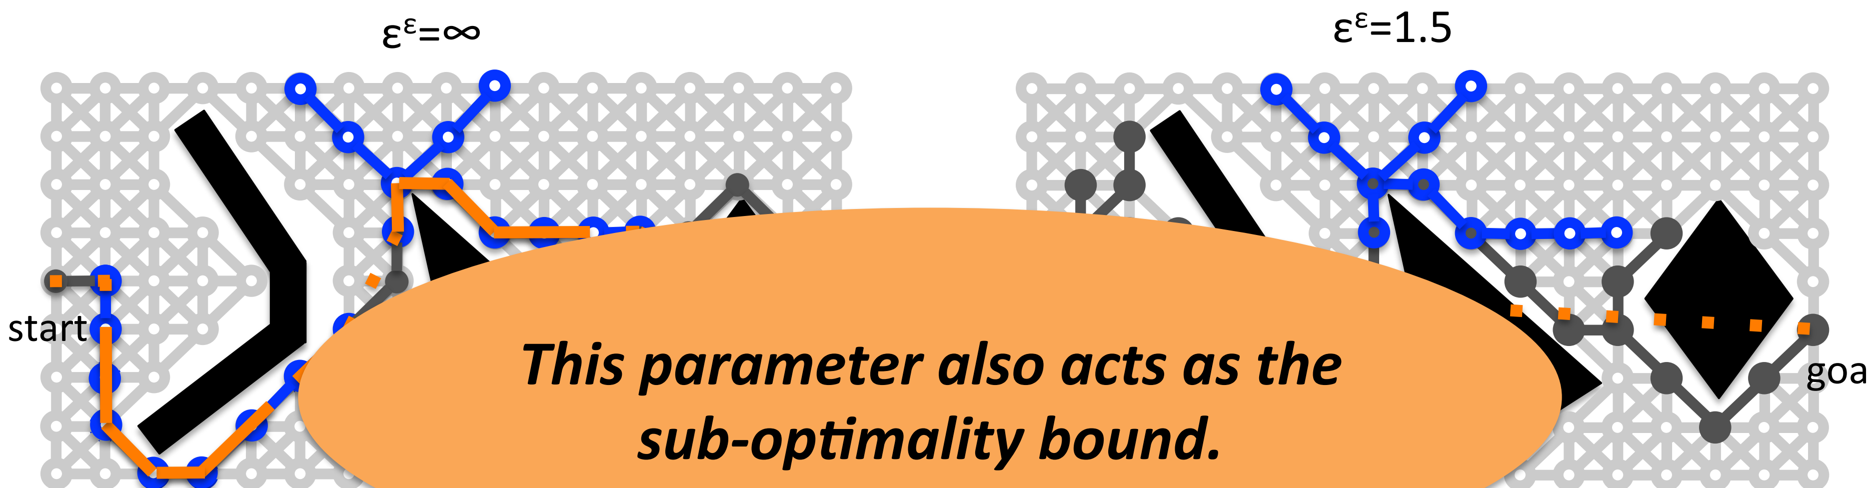

This parameter also acts as the sub-optimality bound.

"E-Graphs: Bootstrapping Planning with Experience Graphs"
Mike Phillips, Benjamin Cohen, Sachin Chitta, Maxim Likhachev
RSS 2012

Experiments on Real and Simulated PR2

- High-Dimensional problems
  - 7 DoF single arm
  - 10 DoF full-body
- Comparison against
  - Weighted A*
  - RRT-Connect
  - PRM
  - RRT*

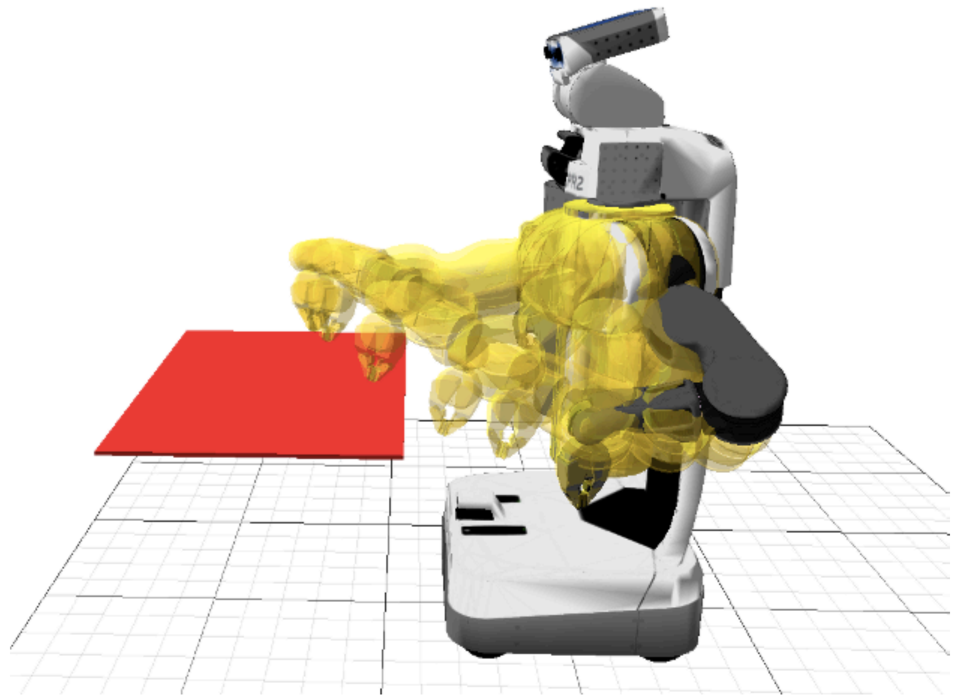

Experiments on Real and Simulated PR2

- High-Dimensional problems
  - 7 DoF single arm
  - 10 DoF full-body
- Comparison against
  - Weighted A*
  - RRT-Connect
  - PRM
  - RRT*
- Results
  - Timing is as fast as sampling methods
  - **Better quality in complex scenarios** where shortcutting is less helpful
  - Much **more consistent** plans

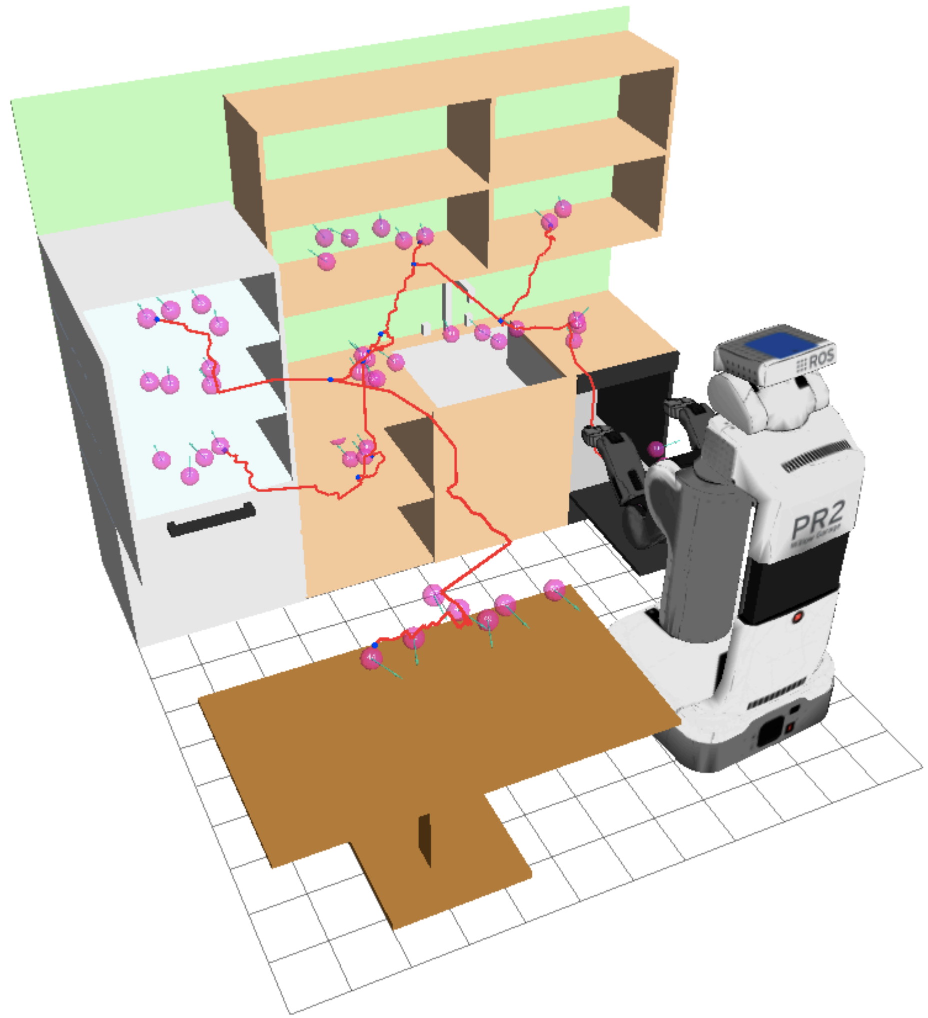

Anytime Planning

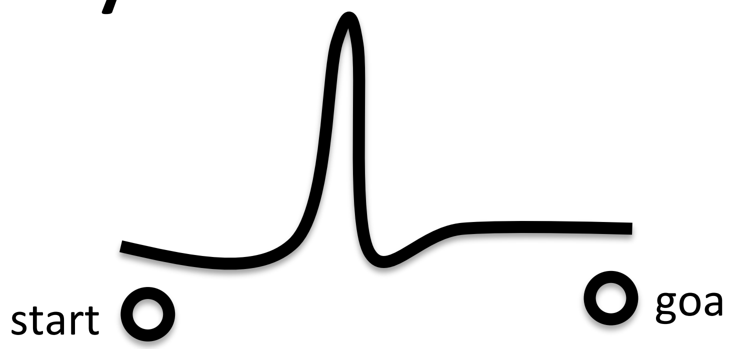

As $\varepsilon^{\mathcal{E}}$ decreases we'd like to see a less dependence on prior experience

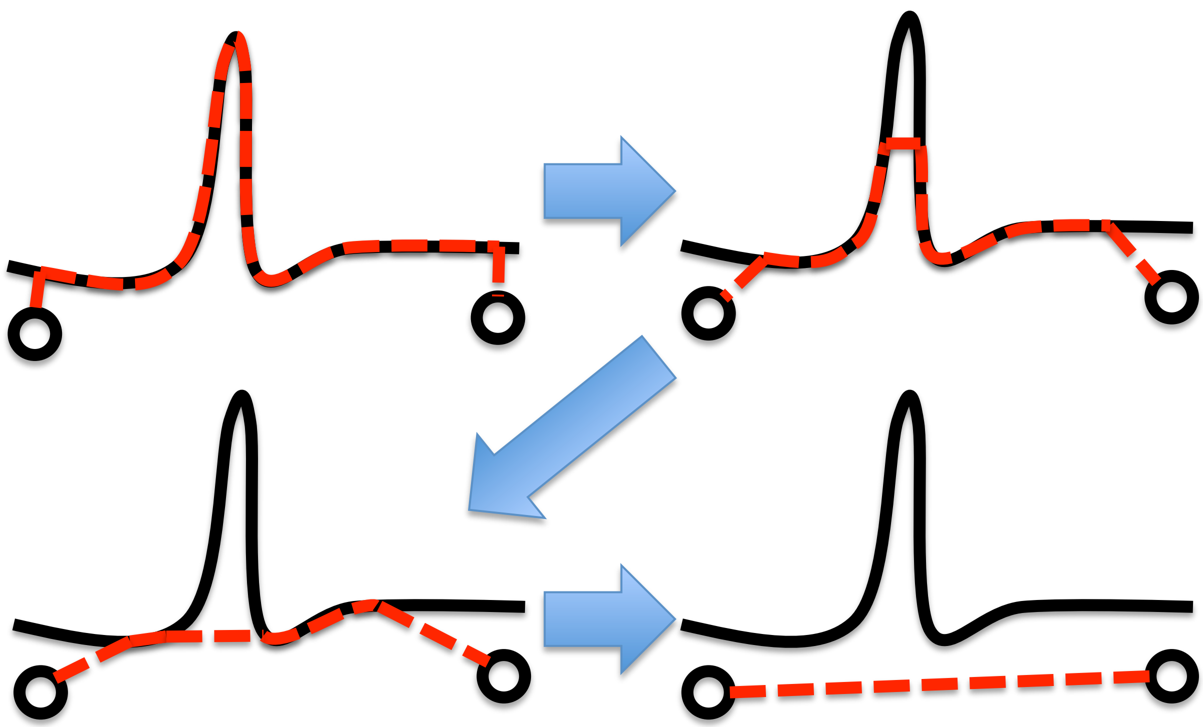

Incremental Planning

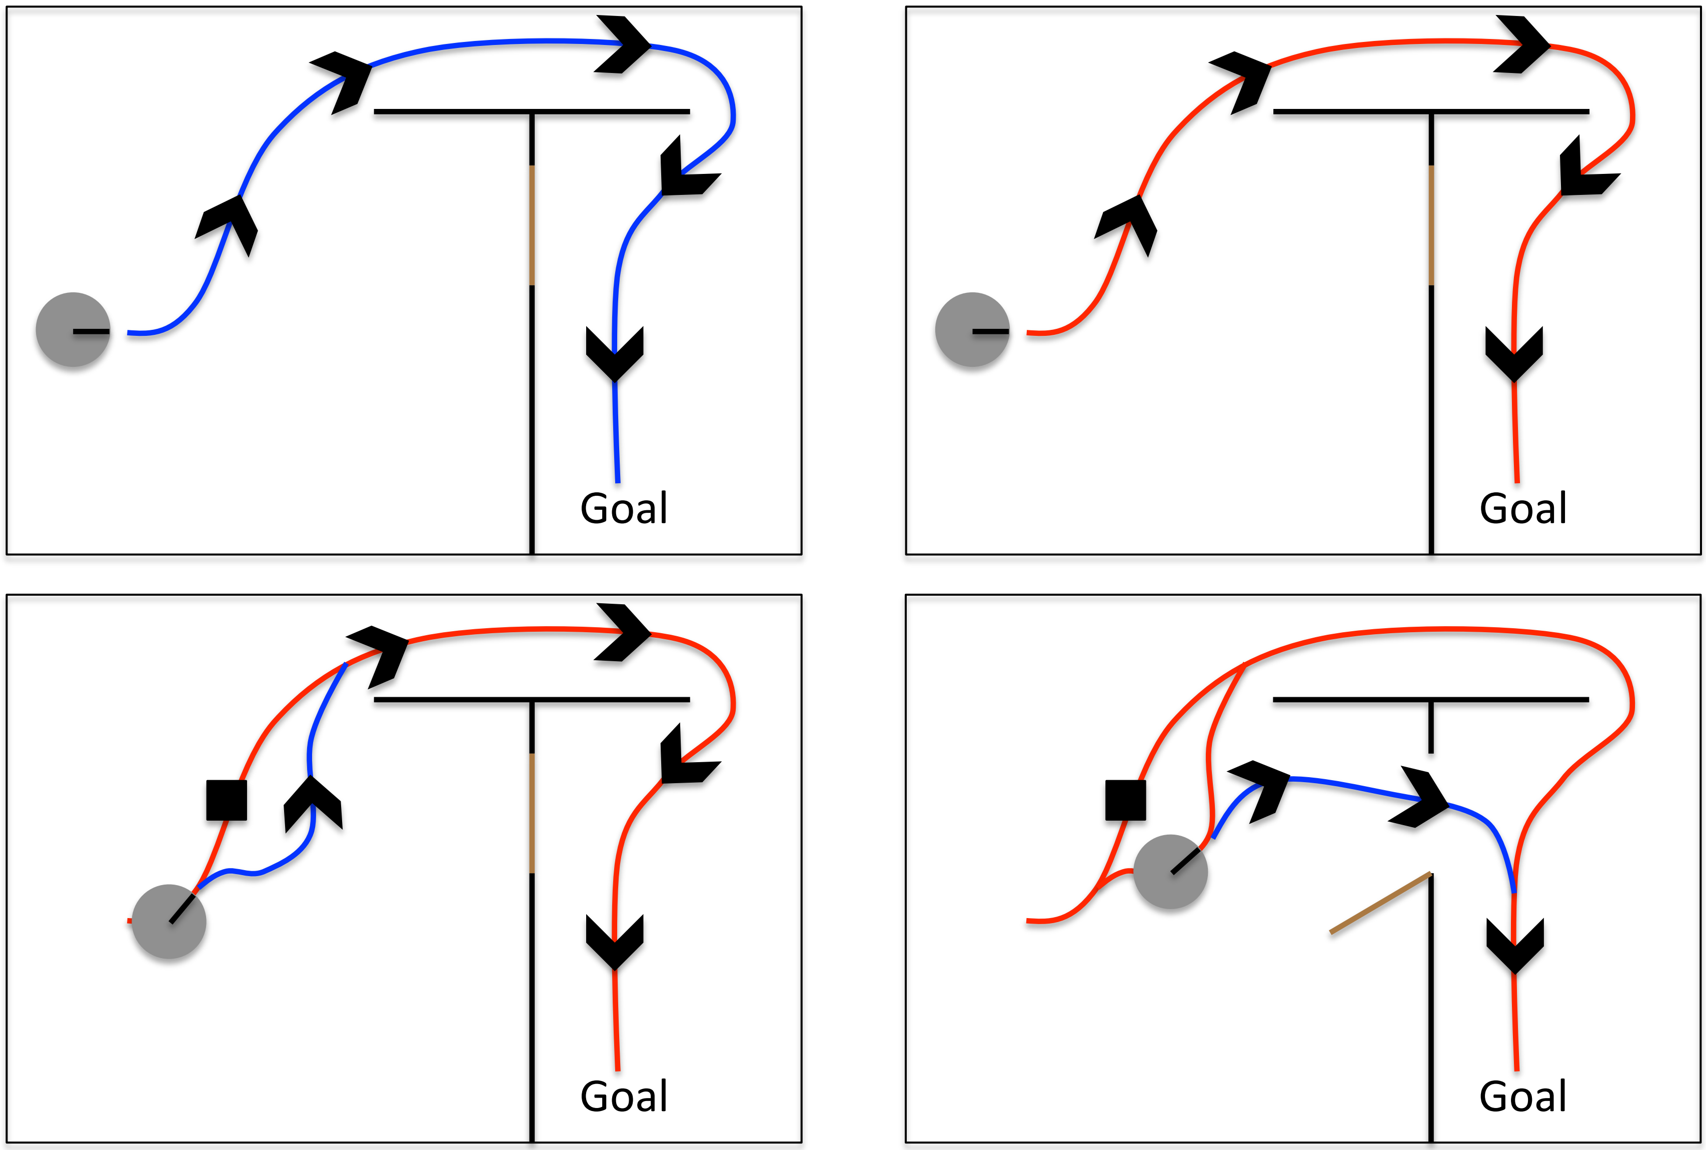

[redacted]

Conclusion

- Experience Graphs use previous plans to accelerate future planning
- Unlike previous approaches, E-Graphs allow for “soft” reuse of parts of experiences
- Theoretical bounds solution cost
- Experiments show planning times on par with sampling methods but better quality and more consistent paths
- Can be used as an anytime planner
- A natural approach to incremental planning