Multi-Agent Path Finding (MAPF) and Its Applications

USC Viterbi School of Engineering

Sven Koenig
skoenig@usc.edu
idm-lab.org

Abstract

Teams of robots must often assign locations among themselves and then plan collision-free paths to these locations. Examples include autonomous aircraft towing vehicles, automated warehouse systems, office robots, drones and game characters in video games. For example, soon, autonomous aircraft towing vehicles will tow aircraft all the way from the runways to their gates (and vice versa) to avoid the pollution and excessive energy consumption resulting from planes taxiing with their engines. Today, hundreds of robots already navigate autonomously in Amazon fulfillment centers to move inventory pods all the way from their storage locations to the inventory stations that need the products they store (and vice versa). Optimal and even some approximately optimal path planning for these robots is NP-hard, yet one must find high-quality collision-free paths for them in real-time since shorter paths result in higher throughput or lower operating costs (since fewer robots are required). Algorithms for such multi-agent path-finding problems have been studied in robotics and theoretical computer science for a longer time but are insufficient since they are either fast but of insufficient solution quality or of good solution quality but too slow (especially for highly congested environments), see my educational website mapf.info for more information.

We therefore study different variants of multi-agent path finding (MAPF) problems (including "lifelong" variants), their complexities, algorithms for solving them and their applications. We proved that even approximating optimal solutions can be NP-hard in some cases but also managed to scale up bounded-suboptimal planners to hundreds of agents and high-quality planners without solution guarantees to thousands of agents via a variety of techniques, such as "highway" heuristics, symmetry-breaking techniques, better node-splitting techniques, mutex propagation, more informed heuristics and learning which collision to resolve next as well as more standard techniques, such as rapid random restarts and large neighborhood search. We also studied MAPF solving based on prioritized planning and deep reinforcement learning. Finally, we developed a hierarchical planning architecture that combines ideas from AI, optimization and robotics. It makes use of a simple temporal network to post-process the output of a MAPF algorithm in polynomial time to create a plan-execution schedule that takes the maximum translational and rotational velocities of robots into account, provides a guaranteed safety distance between them and exploits slack to absorb imperfect plan executions and avoid time-intensive re-planning in many cases.

Science of Making Good Decisions

- We need a science of making good decisions
  - Artificial intelligence
    - Search, planning, coordination, probabilistic reasoning, constraint reasoning, machine learning, …
  - Decision and utility theory
    - utility theory, limited rationality, …
  - Economics
    - auctions, …
  - Operations research
    - Markov decision processes, …
  - Control theory
  - Theoretical computer science and mathematics

The future of AI is in combining decision-making techniques from different disciplines.

Science of Making Good Decisions

Multi-Agent Path Finding (MAPF)

- Optimization problem with the objective to minimize the task-completion time (called makespan) or the sum of travel times (called flowtime)

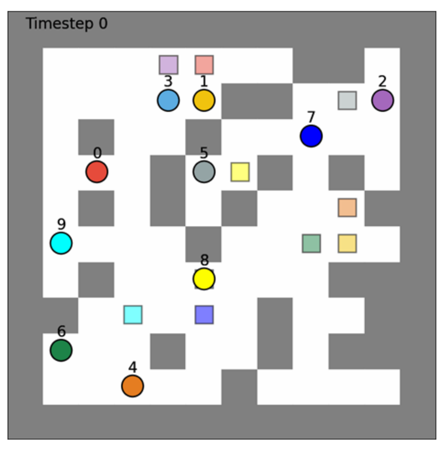

Simplifying assumptions

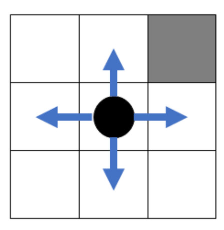

Automated Warehousing

not our work

- Amazon fulfillment centers – warehousing part

https://www.machinedesign.com/mechanical-motion-systems/article/21835788/changing-the-future-of-warehouses-with-amazon-robots

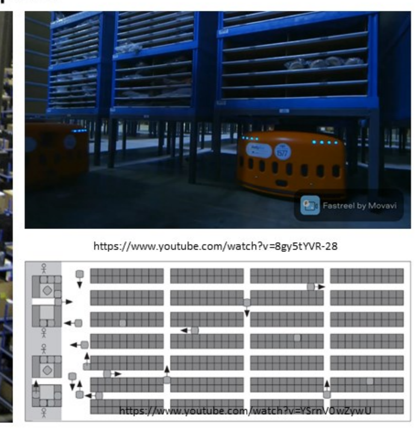

https://www.youtube.com/watch?v=r8gv5tYVR-28

[1] P. Wurman et al., "Coordinating Hundreds of Cooperative, Autonomous Vehicles in Warehouses", AI Magazine 29(1), 9-20, 2008.

Automated Warehousing

not our work

- Amazon fulfillment centers – sorting part

https://www.wired.com/story/amazon-warehouse-robots/

Automated Warehousing

- 800 robots (= 32% empty cells) on a 37x77 sorting-center map with 50 workstations and 275 chutes (joint project with Amazon Robotics)

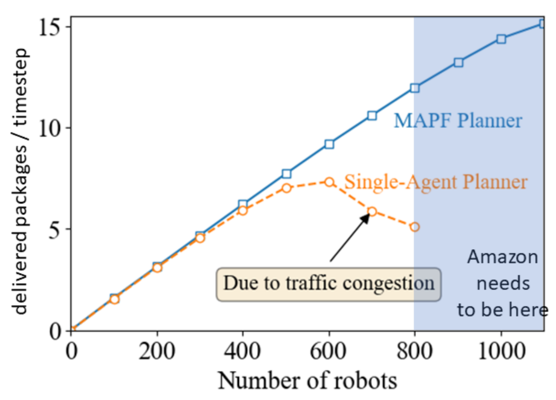

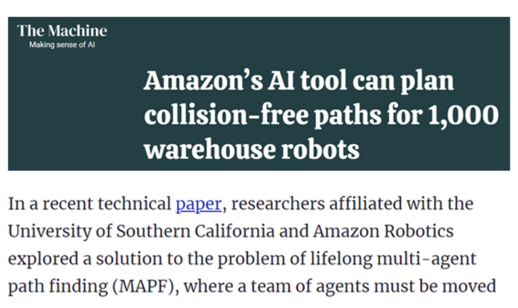
The Machine
Amazon's AI tool can plan collision-free paths for 1,000 warehouse robots

In a recent technical paper, researchers affiliated with the University of Southern California and Amazon Robotics explored a solution to the problem of lifelong multi-agent path finding (MAPF), where a team of agents must be moved

https://venturebeat.com/2020/05/18/amazons-ai-tool-can-plan-collision-free-paths-for-1000-warehouse-robots/

[2] J. Li et al., "Lifelong Multi-Agent Path Finding in Large-Scale Warehouses", AAAI, 2021.

Complexity of MAPF

- Optimal or bounded-suboptimal MAPF planning
  - NP hard to find a makespan- or flow-time optimal MAPF plan [3], even on planar [4] or grid-like graphs [5]
  - NP-hard to find a makespan-bounded-suboptimal MAPF plan with suboptimality factors of less than 4/3 [6]

https://www.pngegg.com/en/png-horua

[3] J. Yu and S. LaValle, "Structure and Intractability of Optimal Multi-Robot Path Planning on Graphs", AAAI, 2013.
[4] J. Yu, "Intractability of Optimal Multi-Robot Path Planning on Planar Graphs", IEEE Robotics and Automation Letters, 2016.
[5] J. Banfi et al., "Intractability of Time-Optimal Multi-Robot Path Planning on 2D Grid Graphs with Holes", IEEE Robotics and Automation Letters, 2017.
[6] H. Ma et al., "Multi-Agent Path Finding with Payload Transfers and the Package-Exchange Robot-Routing Problem", AAAI, 2016.

Planning-Based MAPF Algorithms

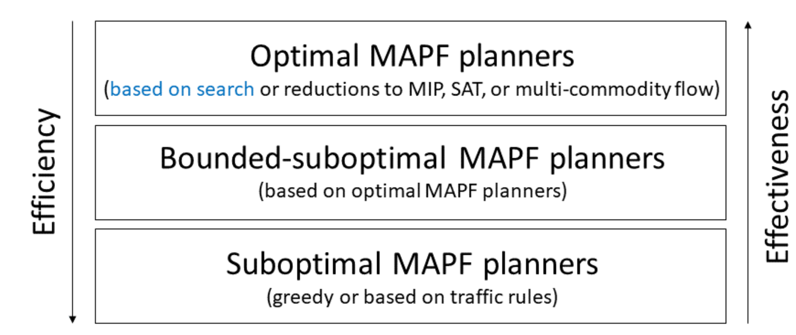

[7] E. Boyarski et al., "Iterative-Deepening Conflict-Based Search", IJCAI 2020.

Planning-Based MAPF Algorithms

- Runtime limit: 1 minute

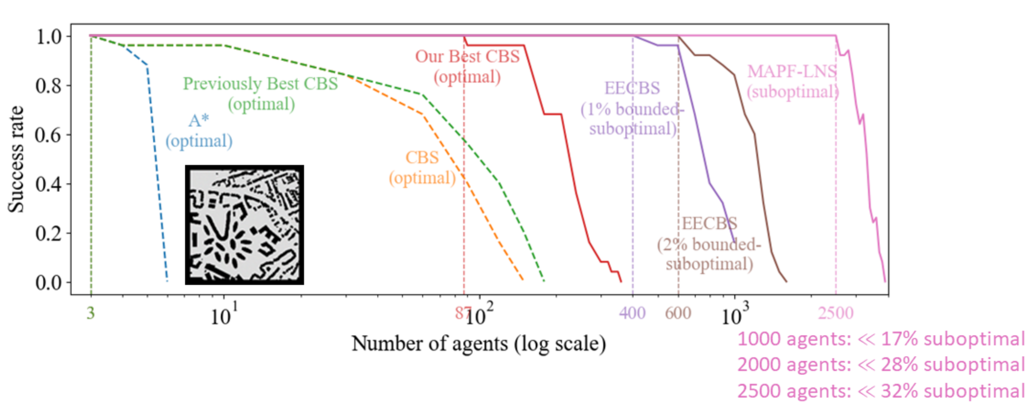

Planning-Based MAPF Algorithms

- Resolving collisions efficiently
  - Symmetry breaking with designed or discovered constraints (optimal)
  - Disjoint splitting (optimal)
  - Priority-based planning (suboptimal)
- Selecting nodes cleverly
  - Informed heuristics (optimal)
  - Explicit Estimation Search (bounded-suboptimal)
- Reducing the number of collisions in a node
  - Highways (bounded-suboptimal)
  - Rolling-horizon collision resolution (suboptimal)
- Other techniques
  - Random restarts (optimal)
  - Large neighborhood search with designed or learned neighborhoods (suboptimal)

Planning-Based MAPF Algorithms

- Agents of more complicated kinodynamics (left)
- Uncertainty about the speeds of agents during execution (right)

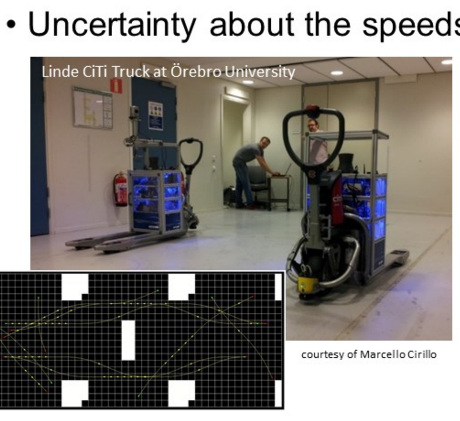

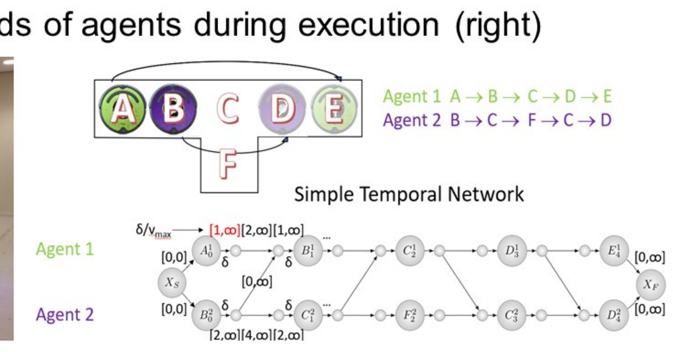

MAPF-POST

[8] L. Cohen et al., "Optimal and Bounded-Suboptimal Multi-Agent Motion Planning". SOCS 2019.
[9] W. Hoenig et al., "Multi-Agent Path Finding with Kinematic Constraints", ICAPS, 2016.

Learning-Based MAPF Algorithms

not our work

- PRIMAL: mix of deep reinforcement learning and imitation learning

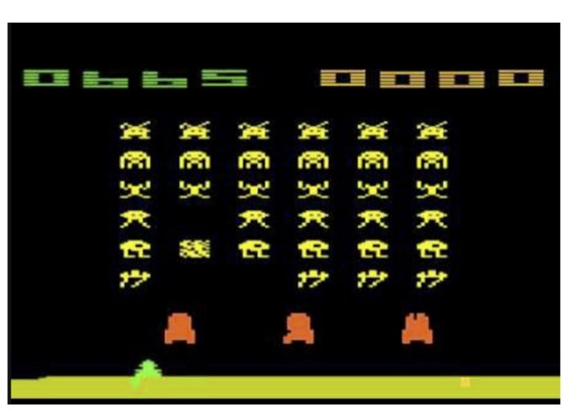

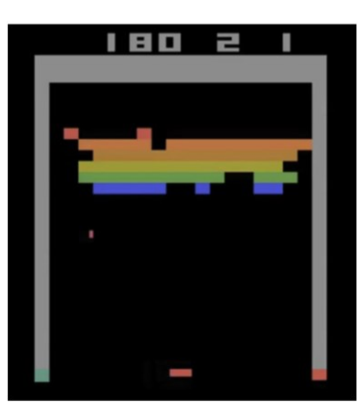

https://medium.com/free-code-camp/explained-simply-how-deepmind-taught-ai-to-play-video-games-9eb5f38c89ee https://www.pcworld.com/article/2889432/google-ai-program-masters-classic-atari-video-games.html

[10] V. Mnih et al., "Human Level Control through Deep Reinforcement Learning", Nature 518, 2015.

Learning-Based MAPF Algorithms

- PRIMAL: mix of deep reinforcement learning and imitation learning

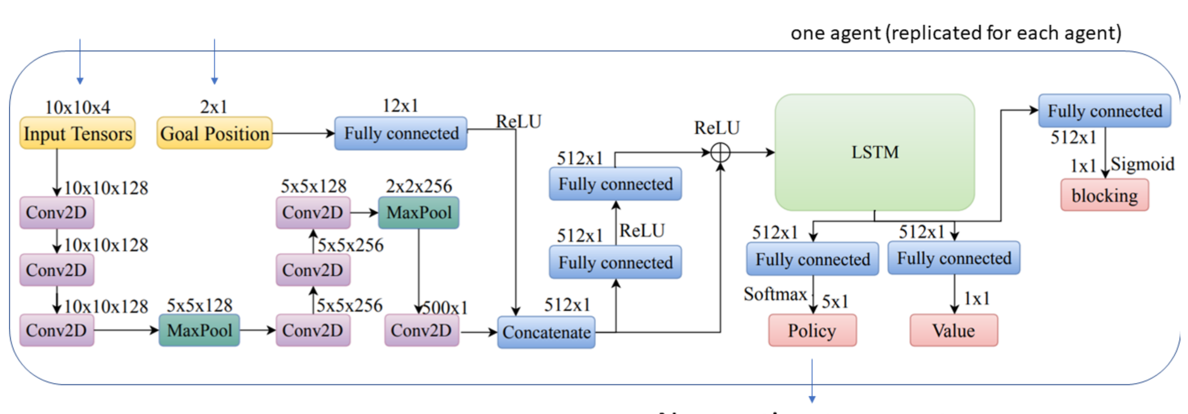

Next action to execute

[11] G. Sartoretti et al., "PRIMAL: Pathfinding via Reinforcement and Imitation Multi-Agent Learning," IEEE RAL 4(3), 2019.

Learning-Based MAPF Algorithms

- PRIMAL: mix of deep reinforcement learning and imitation learning

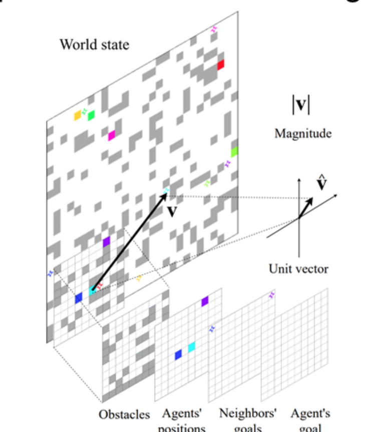

[12] G. Sartoretti et al., "PRIMAL: Pathfinding via Reinforcement and Imitation Multi-Agent Learning," IEEE RAL 4(3), 2019.

Learning-Based MAPF Algorithms

- Advantages
  - Learning takes time but following the resulting policy is fast during execution
  - Can yield a decentralized/distributed system while most effective planning based MAPF algorithms are centralized
- Challenge
  - End-to-end learning does not yet work sufficiently well, as also shown by the results of the NeurIPS 2020 Flatland competition

Learning-Based MAPF Algorithms

- NeurIPS 2020 Flatland Competition: 700 participants from 51 countries

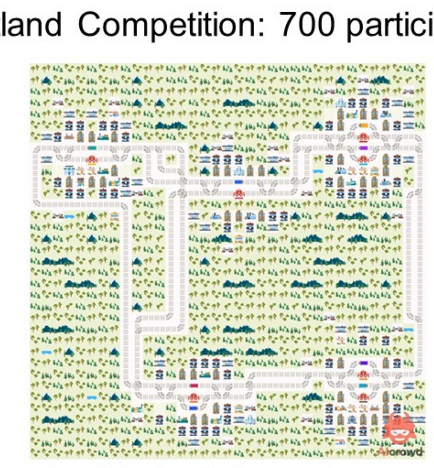

thousands of trains in large environments

[13] J. Li et al., "Scalable Rail Planning and Replanning: Winning the 2020 Flatland Challenge," ICAPS, 2021.

Learning-Based MAPF Algorithms

- NeurIPS 2020 Flatland Competition: 700 participants from 51 countries

Team of
USC and Monash University

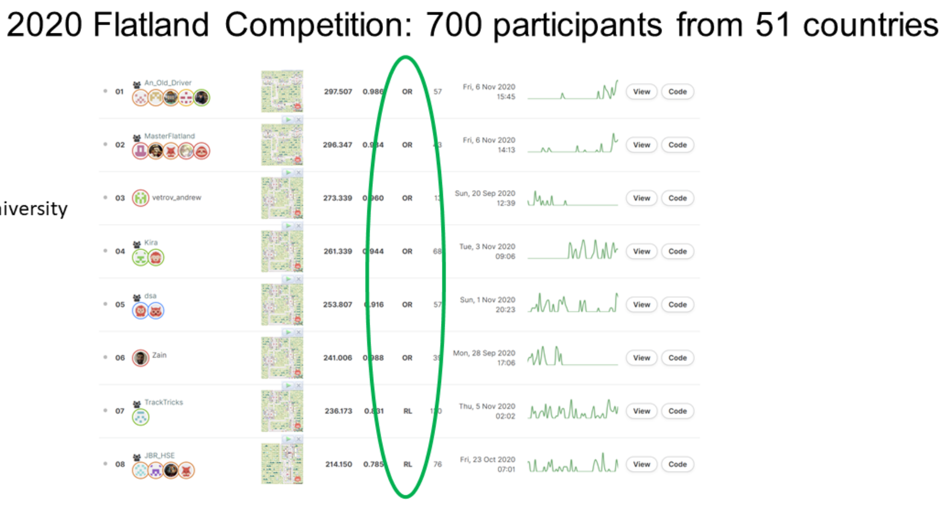

[14] J. Li et al., "Scalable Rail Planning and Replanning: Winning the 2020 Flatland Challenge," ICAPS, 2021.

Learning-Based MAPF Algorithms

- NeurIPS 2020 Flatland Competition: 700 participants from 51 countries

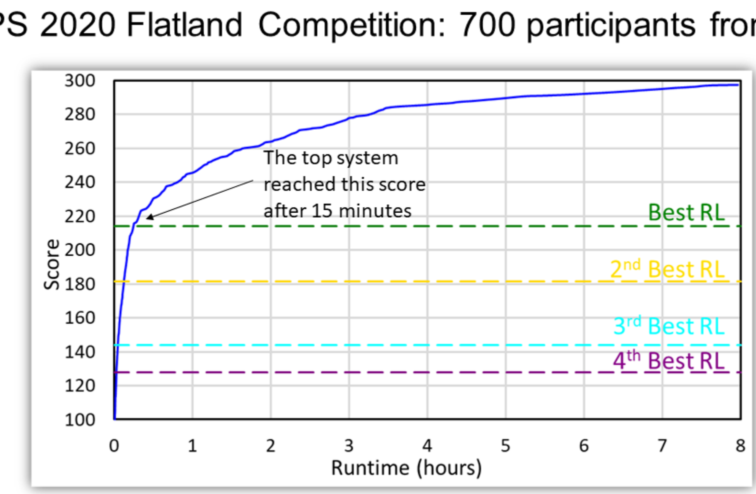

Score of the overall top system

Score of the top learning-based system

[15] J. Li et al., "Scalable Rail Planning and Replanning: Winning the 2020 Flatland Challenge," ICAPS, 2021.

Learning-Based MAPF Algorithms

- Enhancing planning with machine learning
  - Planners use lots of hard-coded decision strategies
  - Machine learning can often learn to make better decisions
  - The resulting planners can be more efficient and/or effective

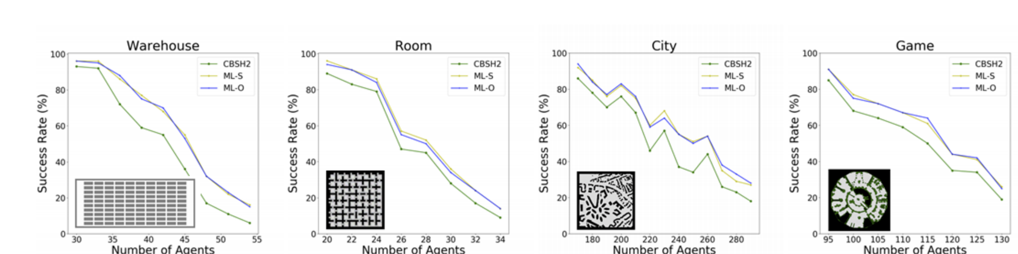

[16] T. Huang et al., "Learning to Resolve Conflicts for Multi-Agent Path Finding with Conflict-Based Search," AAAI, 2021.
[17] T. Huang et al., "Learning Node-Selection Strategies in Bounded-Suboptimal Conflict-Based Search for Multi-Agent Path Finding", AAMAS, 2021.

Multi-Arm Assembly

[18] J. Chen et al., "Cooperative Task and Motion Planning for Multi-Arm Assembly Systems", under review for IEEE Robotics and Automation Letters and ICRA, 2022.

Airport Surface Operation

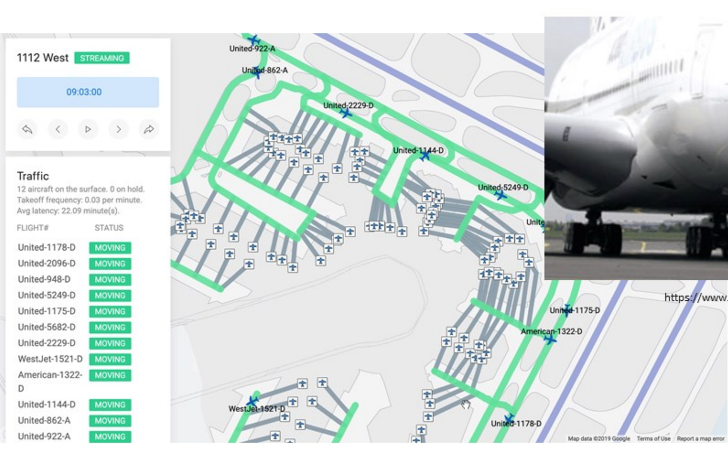

https://www.kessler-axles.de/en/markets/aircraft-ground-support/aircraft-towing-tractors/

[1] J. Li et al., "Scheduling and Airport Taxiway Path Planning under Uncertainty", AIAA, 2019.
[2] R. Morris et al., "Planning, Scheduling and Monitoring for Airport Surface Operations", AAAI-16 Workshop, 2016.

Multi-Agent Path Finding (MAPF)

- Want to learn more about multi-agent path finding?
  - Go here: http://mapf.info/

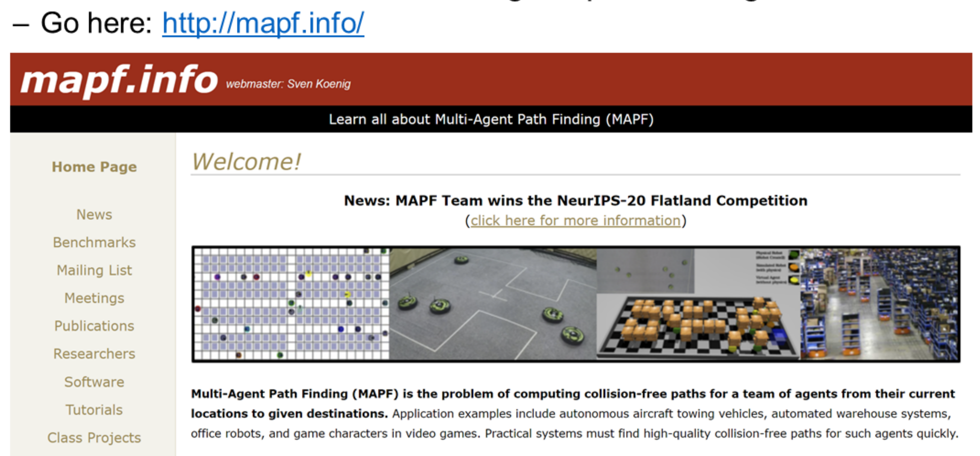

Deserving All Credit

- H. Andreasson, D. Atzmon, N. Ayanian, J. Baier, R. Bartak, G. Belov, E. Boyarski, V. Bulitko, T. Cai, D. Chan, S.-H. Chan, J. Chen, Z. Chen, H. Choset, M. Cirillo, L. Cohen, B. Dilkina, W. Du, J. Durham, A. Felner, G. Gange, M. Garcia de la Banda, M. Gong, M. Greco, D. Harabor, C. Hernandez, W. Hoenig, T. Huang, S. Jung, J. Kerr, S. Kiesel, S. Kumar, E. Lam, P. Le Bodic, J. Li, Z. Liang, M. Liu, W. Liu, Z. Liu, H. Ma, R. Morris, W. Paivine, C. Pasareanu, F. Pecora, W. Ruml, G. Sartoretti, G. Sharon, Y. Shi, D. Sigurdson, R. Stern, P. Stuckey, N. Sturtevant, K. Sun, P. Surynek, L. Terr, A. Tinka, Z. Tong, C. Tovey, T. Uras, G. Wagner, T. Walker, J. Wang, X. Wei, B. Williams, Y. Wu, H. Xu, M. Yao, W. Yeoh, L. Yi, D. Zhang, H. Zhang, Y. Zheng and many others (my apologies that I ran out of space)
- We thank NSF, Amazon Robotics, and NASA Ames for funding several of our MAPF projects!