Hang Ma

Research Statement

Artificial intelligence (AI) is the key to the next wave of industrial revolution. We are stepping into a future where teams of service robots, industrial robots, self-driving cars, autonomous drones, and other robots will become an intrinsic part of our daily life. Having worked on multidisciplinary research topics, I observe that AI has reached a critical point with many new applications and exciting future directions. These future directions, e.g., building teams of robots and other intelligent agents that benefit humanity, rely on a proper blend of advancements in AI and other research areas as well.

Therefore, **my vision** is to bridge the gap between the core AI research and research in other areas and push the resulting interdisciplinary research in a direction to tackle fundamental computational challenges in the real world. **My research goal** is to develop a methodology to provide algorithmic solutions for making intelligent decisions in these challenging real-world scenarios for the benefit of humanity so as to, e.g., automate labor-intensive or dangerous jobs, improve environmental sustainability, and promote social good. To achieve this goal, I unify tools and techniques from the AI, robotics, operations research, and theoretical computer science communities. So far, I have established **an algorithmic foundation for making fast and good decisions to coordinate the long-term operations for real-world multi-agent systems** at a scale of hundreds of robots and thousands of tasks [1, 2, 3]. My contributions to tackling these computational challenges have resulted from applying my expertise in automated planning, combinatorial search, multi-robot systems, combinatorial optimization, spatio-temporal and constraint reasoning, and graph and complexity theory. In the coming years, I will strive to bring together researchers from more areas within and outside computer science including civil, mechanical, and electrical engineering, economics, philosophy, physics, and sociology to help teams of robots, other intelligent agents, and even humans to make good decisions quickly.

Current Research

My current work brings AI and robotics closely together. It opens up many research directions for coordinating multiple agents in real-world autonomous systems. In these systems, teams of agents must constantly assign tasks among themselves and then plan collision-free paths to the task locations. Examples include autonomous aircraft-towing vehicles, automated warehouse robots, automated guided port vehicles, search-and-rescue robots, swarms of ground and aerial robots, game characters in video games, robots in formations, and other multi-robot systems. For example, together with researchers from NASA's Ames research center, I envision that, in the coming years, autonomous aircraft-towing vehicles will tow aircraft all the way from the runways to the gates assigned to them (and vice versa), thereby reducing pollution, energy consumption, congestion, and human workload [4]. With support from Amazon Robotics and Alibaba, I work on coordination problems where thousands of robots navigate autonomously in warehouses to move shelves of products all the way from their storage locations to the packing stations that need the products they contain (and vice versa). Coordinating the agents in these systems is computationally challenging, yet one must find high-quality robust solutions in real time. Solutions that are more efficient, effective, and robust can result in higher throughput, lower operating costs (since fewer robots are required), and more robustness for the long-term autonomy of the whole system.

Thesis Research

My thesis work addresses the following **central question**: *How can one make fast and good decisions to coordinate the long-term autonomous task- and path-planning operations for the above applications of multi-agent systems?*

I formalize and study novel combinatorial problems, e.g., **Combined Task Assignment and Path Finding (TAPF)** [2] and **Multi-agent Pickup and Delivery (MAPD)** [3], that capture both the task- and path-planning aspects of the coordination problems in these multi-agent systems. I then establish an algorithmic framework that exploits the combinatorial structure of the problems, allows for efficient and effective solutions for the whole system, and utilizes domain-specific environmental characteristics to ensure the robustness for the long-term autonomy of these systems. As an example, this framework scales to 250 robots and 2,000 tasks (Figure 1) using 10 seconds of runtime on a laptop computer. It results from three phases of my thesis work, which leverage insights and tools from (1) operations research and theoretical computer science to characterize the computational difficulty and capture the combinatorial structure, (2) AI to provide algorithmic solutions that are efficient and effective by exploiting the combinatorial structure, and (3) robotics to ensure system robustness:

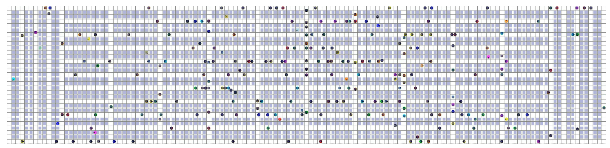

Figure 1. Simulated automated warehouse with 250 robots.

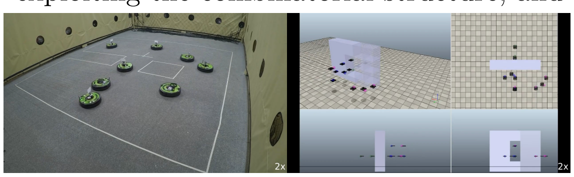

Figure 2. Demos of CBM on ground and aerial robots.

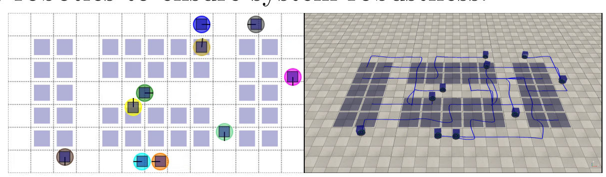

Figure 3. Robots executing kinematically feasible schedules.

1. I formalize TAPF that incorporates task-planning operations, e.g., assigning task locations to teams of robots, into the framework of Multi-Agent Path Finding [5], that was originally designed for path-planning operations only. I conduct a theoretical study and prove the NP-hardness of approximating the optimal TAPF solution within a small constant factor [1]. The results of this study lay the theoretical foundation for many problems I work on, including multi-robot item transportation with package transfers [1] and multi-agent path planning for tasks with deadlines [6, 7] and, as a side benefit, provide insights into many other research problems (e.g., improving the state-of-the-art complexity results for Multi-Agent Path Finding). My work in this phase answers the question of *how hard it is to coordinate the task- and path-planning operations for the target applications.*
2. I develop a hybrid algorithm, called Conflict-Based Min-Cost Flow (CBM), that is complete and optimal for the NP-hard problem of TAPF [2]. CBM exploits the combinatorial structure of TAPF by breaking it down to the NP-hard sub-problem of coordinating different teams of agents and the polynomial-time solvable sub-problems of coordinating agents in every team. It then tackles these sub-problems by using a combination of a combinatorial search algorithm for the NP-hard sub-problem and a min-cost max-flow algorithm for the polynomial-time solvable sub-problems. It scales to more than 400 agents within one minute of runtime in a simulated automated warehouse. Together with robotics researchers, I also apply it to swarms of ground and aerial robots (Figure 2) [8, 9]. My work in this phase answers the question of *how to (and how well one can) coordinate the task- and path-planning operations for the target applications.*
3. I formalize MAPD to allow for robustness in the long-term task- and path-planning operations.

I develop two decoupled algorithms, TP and TPTS, that are extendable to fully distributed settings, and a centralized algorithm, CENTRAL, that is more effective [3]. These algorithms use domain-specific environmental characteristics to bridge the gap between the task- and path-planning components and add to the robustness of the system because they avoid deadlocks that can result from existing task- and path-planning frameworks. I further improve their efficiency and effectiveness by using a novel data structure for collision checking. The resulting algorithm TP-SIPPwRT takes kinematic constraints of real robots into account directly during planning, computes continuous schedules that work on non-holonomic robots, and guarantees a safety distance between the robots (Figure 3) [10]. My work in this phase fully answers the central question: *How can one make fast and good decisions to coordinate the long-term autonomous task- and path-planning operations for the above applications of multi-agent systems?*

Other Research

I also lead several projects that broaden the impact of my research to address other concerns in real-world multi-agent/robot systems.

- For systems with uncertain kinematics and dynamics, I develop a planning and execution algorithmic framework for robust plan execution under uncertainty [11]. It uses a probabilistic model of uncertainty in agent motions and robust plan-execution policies, that together bridge the gap between planning algorithms and the imperfection in plan execution by simultaneously reasoning about precedence/causal constraints required for task-level coordination and temporal constraints required for motion-level coordination. It also avoids intensive communication between robots when used in a distributed setting.
- For tasks with temporal constraints, I formalize the problem of Multi-Agent Path Finding with Deadlines and study the complexity of solving it optimally. Despite the NP-hardness of the problem, I develop several optimal algorithms that scale to dozens of robots in minutes and can eventually be used in real-world applications, e.g., to maximize profits for robots meeting deadlines of the shipping orders in an automated warehouse or to minimize loss for robots evacuating before a disaster occurs in inclement or adversarial conditions in an extraterrestrial exploration [7].
- For large-scale systems that require real-time path-planning computation of high-quality solutions, I develop a novel prioritized path-planning framework for multi-agent systems [12] that scales to 600 agents in 30 seconds of runtime and produces close-to-optimal solutions.
- In addition to existing real-world multi-robot systems, I also generalize my framework to many different application domains, e.g., planning for video game characters [13] and other domains that involve decision making for multiple agents, as described in a survey paper [14].

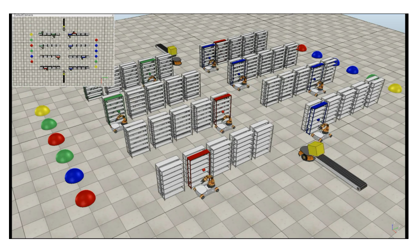

Figure 4. Case study: Plan generation and execution for industrial robots.

Ever since I established this algorithmic foundation for these coordination problems, I have also strived to excite and bring together researchers from different research areas around the world to work on problems that can further broaden the impact of my research. As part of these interdisciplinary and international teams, I have contributed many algorithmic techniques with my expertise in different research areas. For example, we have found great success in improving the efficiency of Multi-Agent Path Finding algorithms by combining combinatorial search techniques with techniques for solving integer linear programming, Boolean satisfiability, and constraint satisfaction problems and leveraging insights from robotics to exploit the problem structure [15, 16, 17]. We have also developed an anytime multi-agent planning algorithm by leveraging incremental search techniques [18]. Finally, together with other robotics researchers, I have developed a hierarchical framework for plan generation and execution in multi-robot systems that brings the above task- and path-planning algorithms one step closer to real-world multi-robot systems. It uses planning algorithms for generating plans that achieve high-level tasks for teams of robots and provides a method for absorbing an imperfect plan execution to avoid time-consuming re-planning in many cases (Figure 4) [8, 19, 9].

Research Agenda

Although teams of robots have become part of our daily life today, many of them are still far from intelligent agents. Therefore, my vision for the next decade is to apply my expertise in different research areas to provide algorithmic solutions for teams of robots, other intelligent agents, and even humans to make intelligent decisions, e.g., to reduce usage of land as a resource, energy consumption, deforestation, pollution, human workload, and safety hazards for humans. So far, I have made progress in real-world multi-agent systems that use hundreds of robots to automate labor-intensive and dangerous jobs. I plan to bring together more researchers from outside AI to build systems that involve thousands or even millions of agents to tackle tasks that require higher-level of intelligence.

My current and projected research address problems with real-world impacts. For example, I was funded by NASA's Ames research center to initiate a blue-sky project on automated airport surface management. I was invited to present my research on task- and path-planning for teams of agents at the ARL West opening ceremony and was later funded to contribute to a project on automated Mars exploration. I was also invited by Cainiao, a logistics company within Alibaba, to their headquarter to give a tutorial talk on long-term planning for large-scale multi-robot systems and have just received funding to initiate a project on coordinating thousands of warehouse robots. My current research is also supported by five NSF grants, a MURI grant, and a gift from Amazon Robotics and has been presented as an essential part of the ACM Turing 50th Celebration conference keynote, several ACM Distinguished Speaker talks, and many other conference and workshop keynotes and tutorials.

I plan to explore new directions to address fundamental questions, including:

- Can my decision-making framework scale to millions of agents by integrating a hierarchical planning approach, which will eventually facilitate autonomous ground and aerial transportation, e.g., for the Los Angeles metropolitan area?
- How can automated planning techniques be combined with data-driven techniques from the robotics, machine learning, and data mining research communities, which will eventually allow for stronger AI, e.g., smarter robots, that can better assist humans?
- How can domain-specific knowledge be integrated into the automated decision-making framework to design environments, e.g., warehouse layouts and road networks, where teams of intelligent agents and humans can collaborate safely, efficiently, and effectively?
- How can AI techniques be helpful to develop specialized algorithmic solutions that are more understandable, explainable, and interpretable and better exploit the problem structure to tackle challenging social problems in sustainability and security domains, e.g., building energy optimization and patrol planning for rangers and air marshals, that currently rely on general black-box solvers, which will eventually reduce energy consumption, prevent and tackle crime, and help build a sustainable environment and human society?

I strive to make impacts that both advance the academic research and improve human society. My interdisciplinary research on large-scale intelligent decision-making frameworks will facilitate the next generation of intelligent agents that make our lives better.

References

[1] **Hang Ma**, Craig Tovey, Guni Sharon, T. K. Satish Kumar, and Sven Koenig. "Multi-Agent Path Finding with Payload Transfers and the Package-Exchange Robot-Routing Problem". *AAAI Conference on Artificial Intelligence (AAAI)*. 2016, pp. 3166–3173.

[2] **Hang Ma** and Sven Koenig. "Optimal Target Assignment and Path Finding for Teams of Agents". *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2016, pp. 1144–1152.

[3] **Hang Ma**, Jiaoyang Li, T. K. Satish Kumar, and Sven Koenig. "Lifelong Multi-Agent Path Finding for Online Pickup and Delivery Tasks". *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2017, pp. 837–845.

[4] Robert Morris, Corina Pasareanu, Kasper Luckow, Waqar Malik, **Hang Ma**, T. K. Satish Kumar, and Sven Koenig. "Planning, Scheduling and Monitoring for Airport Surface Operations". *AAAI-16 Workshop on Planning for Hybrid Systems (PlanHS)*. 2016.

[5] **Hang Ma** and Sven Koenig. "AI Buzzwords Explained: Multi-Agent Path Finding (MAPF)". *AI Matters* 3(3). 2017, pp. 15–19.

[6] **Hang Ma**, Glenn Wagner, Ariel Felner, Jiaoyang Li, T. K. Satish Kumar, and Sven Koenig. "Multi-Agent Path Finding with Deadlines: Preliminary Results [Extended Abstract]". *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2018, pp. 2004–2006.

[7] **Hang Ma**, Glenn Wagner, Ariel Felner, Jiaoyang Li, T. K. Satish Kumar, and Sven Koenig. "Multi-Agent Path Finding with Deadlines". *International Joint Conference on Artificial Intelligence (IJCAI)*. 2018, pp. 417–423.

[8] **Hang Ma**, Wolfgang Hönig, Liron Cohen, Tansel Uras, Hong Xu, T. K. Satish Kumar, Nora Ayanian, and Sven Koenig. "Overview: A Hierarchical Framework for Plan Generation and Execution in Multi-Robot Systems". *IEEE Intelligent Systems* 32(6). 2017, pp. 6–12.

[9] Wolfgang Hönig, T. K. Satish Kumar, **Hang Ma**, Nora Ayanian, and Sven Koenig. "Formation Change for Robot Groups in Occluded Environments". *IEEE/RSJ International Conference on Intelligent Robots and System (IROS)*. 2016, pp. 4836–4842.

[10] **Hang Ma**, Wolfgang Hönig, T. K. Satish Kumar, Nora Ayanian, and Sven Koenig. "Lifelong Path Planning with Kinematic Constraints for Multi-Agent Pickup and Delivery". *AAAI Conference on Artificial Intelligence (AAAI)*. 2019.

[11] **Hang Ma**, T. K. Satish Kumar, and Sven Koenig. "Multi-Agent Path Finding with Delay Probabilities". *AAAI Conference on Artificial Intelligence (AAAI)*. 2017, pp. 3605–3612.

[12] **Hang Ma**, Daniel Harabor, Peter J. Stuckey, Jiaoyang Li, and Sven Koenig. "Searching with Consistent Prioritization for Multi-Agent Path Finding". *AAAI Conference on Artificial Intelligence (AAAI)*. 2019.

[13] **Hang Ma**, Jingxing Yang, Liron Cohen, T. K. Satish Kumar, and Sven Koenig. "Feasibility Study: Moving Non-Homogeneous Teams in Congested Video Game Environments [Demonstration Abstract]". *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*. 2017, pp. 270–272.

[14] **Hang Ma**, Sven Koenig, Nora Ayanian, Liron Cohen, Wolfgang Hönig, T. K. Satish Kumar, Tansel Uras, Hong Xu, C. Tovey, and G. Sharon. "Overview: Generalizations of Multi-Agent Path Finding to Real-World Scenarios". *IJCAI-16 Workshop on Multi-Agent Path Finding (WOMPF)*. 2016.

[15] Ariel Felner, Jiaoyang Li, Eli Boyarski, **Hang Ma**, Liron Cohen, T. K. Satish Kumar, and Sven Keonig. "Adding Heuristics to Conflict-Based Search for Multi-Agent Pathfinding". *International Conference on Automated Planning and Scheduling (ICAPS)*. 2018, pp. 83–87.

[16] Jiaoyang Li, Pavel Surynek, Ariel Felner, **Hang Ma**, T. K. Satish Kumar, and Sven Koenig. "Multi-Agent Path Finding for Large Agents". *AAAI Conference on Artificial Intelligence (AAAI)*. 2019.

[17] Jiaoyang Li, Daniel Harabor, Peter J. Stuckey, **Hang Ma**, and Sven Koenig. "Symmetry-Breaking Constraints for Grid-Based Multi-Agent Path Finding". *AAAI Conference on Artificial Intelligence (AAAI)*. 2019.

[18] Liron Cohen, Matías Greco, **Hang Ma**, Carlos Hernandez, Ariel Felner, T. K. Satish Kumar, and Sven Koenig. "Anytime Focal Search with Applications". *International Joint Conference on Artificial Intelligence (IJCAI)*. 2018, pp. 1434–1441.

[19] Wolfgang Hönig, T. K. Satish Kumar, Liron Cohen, **Hang Ma**, Hong Xu, Nora Ayanian, and Sven Koenig. "Multi-Agent Path Finding with Kinematic Constraints". *International Conference on Automated Planning and Scheduling (ICAPS)*. 2016, pp. 477–485.