Planning for Flexible Human-Robot Co-navigation in Dynamic Manufacturing Environments

James C. Boerkoel Jr. and Julie A. Shah
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
{boerkoel,julie_a_shah}@csail.mit.edu

I. Introduction

An emerging trend in advanced manufacturing is the increased prevalence of mobile, right-sized robots that work safely and flexibly alongside humans across a wide variety of tasks. Currently in such environments, human workers are responsible for completing a variety of assembly tasks, such as welding or torquing parts into place. While the assembly tasks may require the skill and dexterity of a human worker, a mobile robotic assistant could help fetch parts and tools, increasing the overall productivity of the system. The rob@work platform (Figure 1) by Fraunhofer IPA is an example of one such mobile robot that is capable of fetching and picking parts on behalf of its human coworkers. A rob@work consists of a mobile base with four independent and individually actuated wheels, and can be equipped with other sensors and actuators, such as a robotic arm, to successfully interact with its environment. A major challenge is that the manufacturing environments in which these types of robots are designed to operate are densely populated, dynamic, and human-oriented. Human workers intuitively shuffle by each other through the narrow corridors that emerge between conveyor belts, various part bins, and tool carts. In most cases, this type of coordination between humans is accomplished with little or no explicit communication: humans quickly perceive and adapt to the intentions and conventions of their coworkers. We coin the term co-navigation to describe this natural and intuitive coordinated dance that human teams perform with ease. This poses an overarching challenge for a mobile robotic assistant: how to safely integrate itself in this intuitively choreographed shuffle of human workers in the presence of a cluttered, dynamic environment. Next, we highlight the particular challenges that await, along with our proposed approaches for addressing them, en route to realizing our goal of effective human-robot co-navigation.

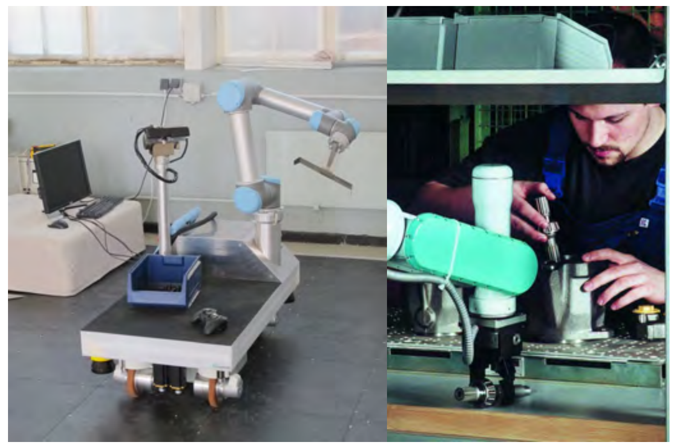

Fig. 1: The rob@work platform by Fraunhofer IPA. Photo Credits: http://idw-online.de and http://care-o-bot-research.org.

II. Open Challenges and Proposed Methods

A. Perceiving and Understanding Human Intentions

Humans successfully collaborate on tasks often without explicit communication by using a variety of implicit social cues and motions to convey intent. We look to exploit many recent advances in human-robot interaction that use implicit human communication to inform plan decisions. For example, features such as the pose and motion of humans have been shown to help identify humans' intended activities [1], while the use of social cues have aided in achieving effective turn-taking with humans [2]. Other recent work uses action observations to infer, and autonomously react to, the intended plan of a human [3]. While our primary contributions will be in the adaptive planning and scheduling of robots in response to their human coworkers, we intend to use these rich features to capture human intent and inform our planning models.

B. Planning Efficacious and Timely Turn-taking

Typically, human teams develop implicit conventions, such as social cues and turn-taking rules, that are learned over time and allow natural negotiation without requiring explicit communication. A primary goal of our work is to equip robots with methods to quickly and robustly develop plans that (1) adapt to established conventions to successfully negotiate human-oriented environments, and (2) augment the workflow of human workers to increase overall productivity, safety, and quality. We will extend previous work that used Partially Observable Markov Decision Processes (POMDPs) to autonomously perform socially situated tasks such as making a left-hand turn or boarding an elevator [3], incorporating recent work in turn-taking [2] and navigation through crowds [4]. We will also build on recent work that learns and exploits a shared mental model between a robot and human-coworker, allowing the robot to flexibly adapt its plan according to the particular workflow patterns of individual human coworkers [5]. We will continue using a POMDP to model and infer (unobservable) intended activities of human coworkers, but we are particularly interested in looking at the interplay between planning which actions to take and the timing of these actions (described in more detail next). We will investigate both which human conventions tend to hold universally, and thus can be explicitly programmed into the robot's planning module, and which elements require on-line learning to adapt to the tendencies of individuals. Next, we discuss using a separate timing representation to fluidly schedule events conditioned on the selected POMDP policies.

C. Low Response-time Robot Dispatching and Scheduling

Minimizing the perceived lag and response time of a system is critical in creating natural, fluid interactions between humans and robots. Unfortunately, explicitly capturing timing in problem state significantly increases planning complexity. Our idea is to continue in the spirit of approaches such as the time-state aggregated POMPD [6] and event-driven interactions [7] to reduce the complexity of planning by using abstraction to decouple the scheduling and timing aspects from the POMDP model. Constraint-based scheduling has proven particularly effective in efficiently dispatching the schedules of multiple interacting robots [8]. Further work has investigated ways to decouple the interdependent schedule of one agent (in our case the robot's) from those of others' (in our case the humans') [9]. This allows the robot to act completely autonomously, thus further decreasing the amount of reasoning required to act. We will extend this work both by exploiting additional problem structure, and by conditioning temporal constraints to be an optimal response to human actions. This leads to non-volitional schedule refinements of two kinds: (i) those to current (or past) events based on humans' actual plan execution; and (ii) those to future events added by the planner in response to perceived human intent. Decoupling timing from planning is not the only abstraction that we expect might yield benefits; explicitly limiting the time horizon has proven also useful in the real-time coordination and control of teams of mobile robots in noisy and dynamic environments [10].

D. Conveying Intended Route

Once a robot has chosen a policy, planned a route, and is ready to schedule the execution of this route, it must effectively communicate its intentions to ensure the continued safety and fluidity of its interactions with a human teammate. There are many modes of implicit or explicit communication that a robot could take. We propose equipping our mobile robots with a laser-based projection system that can both indicate the robot's target destination and also convey its intended near-term route (in terms of heading and velocity) through a projection onto the factory floor. Our hypothesis is that projection, coupled with a robot that can robustly recognize and plan around human teammates' implicit social cues (e.g., pose or facial expressions), will increase the overall timeliness and fluidity of the robot's interactions with humans. We will evaluate this using both quantitative and qualitative metrics, including the overall time required to complete a human-robot shuffling of positions, human coworker's Likert-scale reports of the quality of exchanges, and the qualitative evaluation of domain experts.

E. Evaluation in Human Environment

We will test our algorithms on a mobile robot deployed in a real industrial manufacturing environment. While we will adopt externally developed methods for addressing many perceptual and control challenges, the changing positions of pick carts and a moving conveyor belt provides us a dynamic, cluttered environment in which to evaluate the robustness and generality of our approach. This will allow us to test our socially-adept, co-navigational planning techniques across paths with a wide-variety of clearances and complexities.

III. Expected Contributions

The overarching goal of our work is to achieve high-quality, human-robot co-navigation through the safe and fluid integration of mobile robots into the negotiated choreography of human workers in the presence of a cluttered, dynamic environment. We expect that working towards this goal will lead to the following contributions: (i) a general planning framework that incorporates a rich space of implicit features based on human poses, social cues, and conventions; (ii) a new human-aware approach to planning and scheduling that explicitly decouples timing from the plan model; (iii) new schedule dispatch techniques that exploit underlying problem structure and are immediately responsive to human actions; (iv) a novel approach for conveying intended paths; and (v) an evaluation of our methods in a real, human-oriented manufacturing environments with various navigational hazards. In the future, we hope to extend and evaluate our approaches across a wider variety of both human-robot interaction and planning/scheduling problem domains.

References

[1] J. Sung, C. Ponce, B. Selman, and A. Saxena, "Unstructured human activity detection from rgbd images," in *Proc. of ICRA 2012*, 2012, pp. 842–849.

[2] C. Breazeal, "Towards sociable robots," *Robotics and Autonomous Systems*, vol. 42, no. 3–4, pp. 167–175, 2002.

[3] F. Broz, I. Nourbakhsh, and R. Simmons, "Designing pomdp models of socially situated tasks," in *Proc. of 2011 RO-MAN*, 2011, pp. 39–46.

[4] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *Proc. of IROS-10*, 2010, pp. 797–803.

[5] S. Nikolaidis and J. A. Shah, in *Proc. of HRI-13*, 2013, p. To appear.

[6] F. Broz, I. Nourbakhsh, and R. Simmons, "Planning for human-robot interaction using time-state aggregated pomdps," in *Proc. of AAAI-08*, 2008, pp. 1339–1344.

[7] R. Becker, S. Zilberstein, and V. Lesser, "Decentralized markov decision processes with event-driven interactions," in *Proc. of AAMAS-04*, 2004, pp. 302–309.

[8] J. A. Shah, P. R. Conrad, and B. C. Williams, "Fast distributed multi-agent plan execution with dynamic task assignment and scheduling," in *Proc. of ICAPS-09*, 2009, pp. 289–296.

[9] J. Boerkoel and E. Durfee, "Distributed algorithms for solving the multiagent temporal decoupling problem," in *Proc. of AAMAS-11*, 2011, pp. 141–148.

[10] S. Gowal and A. Martinoli, "Real-time optimization of trajectories that guarantee the rendezvous of mobile robots," in *Proc. of IROS-12*, 2012, pp. 3518 –3525.