MR.CAP: Multi-robot Joint Control and Planning for Object Transport

Hussein Ali Jaafar, Cheng-Hao Kao, and Sajad Saeedi

Abstract—With the recent influx in demand for multi-robot systems throughout industry and academia, there is an increasing need for faster, robust, and generalizable path planning algorithms. Similarly, given the inherent connection between control algorithms and multi-robot path planners, there is in turn an increased demand for fast, efficient, and robust controllers. We propose a scalable joint path planning and control algorithm for multi-robot systems with constrained behaviours based on factor graph optimization. We demonstrate our algorithm on a series of hardware and simulated experiments. Our algorithm is consistently able to recover from disturbances and avoid obstacles while outperforming state-of-the-art methods in optimization time, path deviation, and inter-robot errors. See the code and supplementary video for experiments[1].

Index Terms—robotics, optimization, cooperative control

I. Introduction

IN recent years, there has been a growing demand for cohesive real-time multi-robot systems across various industries, including agriculture, warehousing, and hospitality [1]. These industries have increasingly relied on robotic systems to enhance workflow efficiency and cost-effectiveness. Key challenges in these multi-robot systems revolve around problems such as path planning, control, and estimation. Both planning and control are inherently forward-looking tasks, making use of the current state to predict information on future states. This similarity drives an attempt to integrate these two into one, which leads to reduced inter-module communication, resulting in faster response times and holistically optimal results [2].

In multi-robot systems, the integration of path planning and control into a cohesive real-time framework presents significant challenges. These difficulties manifest as communication overheads, computational inefficiencies, and scalability limits with respect to the number of obstacles and robots. Problem formulation and optimization is the key factor distinguishing the existing solutions [3]. These solutions vary depending on the problem specifics and the required quality. For collaborative object transport [4], hierarchical quadratic programming (HQP) [5] performs partial joint planning and control, where the planning is based only on the current view of the sensor, unlike our work, which takes into account the full map of the environment. Further, scalability and handling uncertainties remain open problems.

To address these challenges, this paper introduces an innovative scalable end-to-end joint planning and control framework. Our approach seamlessly integrates planning and control into a unified optimization framework, harnessing the versatility
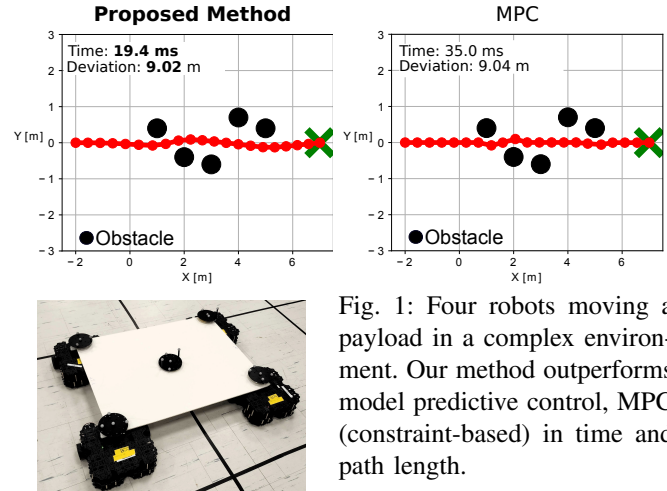
Fig. 1: Four robots moving a payload in a complex environment. Our method outperforms model predictive control, MPC (constraint-based) in time and path length.

of factor graphs [6]. The utilization of factor graphs, with its probabilistic structure, equips our system to handle a wide spectrum of scenarios. Notably, our approach prioritizes scalability, making it easy to accommodate larger robot teams and complex environments without introducing substantial computational complexity. We rigorously evaluate our algorithm against current state-of-the-art approaches, including model predictive control (MPC) [7], [8] and HQP [5], with simulated and hardware-based experiments. The formulation and experiments are designed for collaborative object transport; however, it is flexibly extendable to a wide range of problems. See the project page for demos, a 3D example, and the code.

We focus on joint control and planning for collaborative object transport. The initial formation is assumed to be known. Fig. 1 shows a demo where four robots transport an object while avoiding obstacles. Our method is faster than the MPC-based methods [7] and generates a shorter path. We present the literature review in Sec. II. Sec. III presents our method, followed by experiments in Sec. IV, and conclusions in Sec. V.

II. Literature Review

Motion planning is a cornerstone of robotic systems. Planners can be categorized into several types: graph-based, e.g. A* [9] and Dijkstra [10]; sampling-based, e.g. RRTs [11] and PRM [12]; and trajectory optimizers, e.g. covariant Hamiltonian optimization for motion planning (CHOMP) [13], T-CHOMP [14], Multigrid CHOMP [15], STOMP [16]. Optimizers are typically probabilistic and use *maximum likelihood*, as in [17], or *maximum a posteriori* (MAP), as in simultaneous trajectory estimation and planning (STEAP) [2]. Some of

these methods have high discretization costs, which can be overcome by methods in [18], [19]. Additionally, to reduce the complexity of the methods, and present trajectories with fewer states, Gaussian process motion planning (GPMP) [20] represents the continuous-time trajectory as a sample from a Gaussian process. In many recent planning algorithms, factor graphs play a key role, e.g. [21] and [22]. Planning in dynamic settings (see ITOMP [23]) and lack of guarantee for optimality and feasibility of paths remain an active research field.

In real-world applications, the collaboration of multiple robots is inevitable. There are several research frontiers in such systems, including cooperative exploration [24], localization [25], [26], formation [27], planning [28]–[30], and control [31]. A widely used controller is model predictive control (MPC) [32]. An active research area in these frontiers is to solve the problems in a distributed/decentralized manner, as done in [33] for control, or in [34] for manipulation.

An extension to the multi-robot control and planning problem is object transport/pushing/manipulation [35], using multiple robots to transport an object such that it undergoes a desired motion or tracks a desired path [36]. Various approaches are used for cooperative object transport. Examples include furniture moving algorithm [37], using virtual structures [38], combining primitive controllers [4], hybrid controller with a receding horizon planner [39], constrained optimization [5], and various other interesting centralized and decentralized algorithms [40], [41], [42], [43]. Data-driven models making use of behavioral systems theory have also been proposed for multi-robot quadrupedal locomotion [44], [45]. Distributed scalable manipulation and planning inspired by separable optimization variable ADMM (SOVA) methods have been investigated in [46], [47]. An open problem is the capacity to handle uncertainties, essential for real-world applications.

III. Proposed Method

We tackle a joint planning and control problem of a multi-robot for object transport (see Fig. 2). The robots collaboratively manipulate the object centroid to keep it along a planned path and avoid obstacles in an online fashion. For obstacle avoidance, we require that the entire payload avoid the obstacles along with the robots. The formation of the robots is assumed to be known and remains fixed.

A. Motion Model

The task of joint planning and control is discretized into N steps, i.e. waypoints. There are a total of J obstacles to be avoided. There are two types of variables: the states $\mathbf{x} = \mathbf{x_{0:N}}$, $\mathbf{x}_n = \{x_n, y_n, \theta_n\}$, $n = 1...N$ (see Fig. 2) and controls $\mathbf{u} = \mathbf{u_{0:N-1}}$, $\mathbf{u}_n = \{v_n, \omega_n\}$. Control efforts and velocities in the global frame $\mathbf{v} = \mathbf{v_{0:N-1}}$, $\mathbf{v}_n = \{\dot{x}_n, \dot{y}_n, \dot{\theta}_n\}$ can easily be related by using kinematic Jacobians.

We use a number of I differential drive robots, which can be arranged in an arbitrary pattern underneath the payload. Let the superscripts c and r_i represent the centroid and the i-th robot in the formation respectively, we have centroid poses $\mathbf{x}^c$, centroid controls $\mathbf{u}^c$, robot poses $\mathbf{x}^{r_i}$, and robot controls $\mathbf{u}^{r_i}$. Assuming small integration intervals and constant velocity


Fig. 2: Example of four robots arranged to move the payload placed at an arbitrary location on the platform. The system is capable of operating with more robots, for heavy payloads. The robots can have any starting orientation when connected, as the optimizer accounts for any needed adjustments.

within each step, the motion model of a robot is obtained by integration via the Runge-Kutta method [21]

$$\mathbf{x}_{n+1}^{r_i} = \mathbf{x}_n^{r_i} + T_s \begin{bmatrix} cos(\theta_n^{r_i} + \frac{\omega_n^{r_i}}{2}) & 0 \\ sin(\theta_n^{r_i} + \frac{\omega_n^{r_i}}{2}) & 0 \\ 0 & 1 \end{bmatrix} \mathbf{u}_n^{r_i}, \qquad (1)$$

$$\forall i = 1, 2, \ldots, I \quad \text{and} \quad \forall n = 1, 2, \ldots, N \ .$$

T_s is the processing time step. The motion model of the centroid takes the same form as Eq. (1). However, due to the nonholonomic constraints imposed by the differential drive robots used and the geometric constraints embedded in the system design, we split the motion of the centroid into two stages; translation and rotation. This eliminates the need for Runge-Kutta approximations and improves the path shape of the system. As such, the motion model of the centroid is:

$$\mathbf{x}_{n+1}^c = \mathbf{x}_n^c + T_s \begin{bmatrix} cos(\theta_n^c) & 0 \\ sin(\theta_n^c) & 0 \\ 0 & 1 \end{bmatrix} \mathbf{u}_n^c \ . \qquad (2)$$

Since there are no actuators on the payload, one could also optimize for the centroid velocity $\mathbf{v}_n^c$ in the world frame to further simplify the problem. This means:

$$\mathbf{x}_{n+1}^c = \mathbf{x}_n^c + T_s \mathbf{v}_n^c \ . \qquad (3)$$

Once the control of the centroid is optimized, the robots pivot until they reach their required orientations $\theta_n^{r_i}$ and then apply the controls $\mathbf{u}_n^{r_i}$:

$$\theta_n^{r_i} = \begin{cases} \theta_n^c, & \text{if centroid translates} \\ \theta_n^c + \psi_i \pm \pi/2, & \text{if centroid rotates} \end{cases}, \qquad (4)$$

$$\mathbf{u}_n^{r_i} = \begin{bmatrix} 1 & l_i \\ 0 & 1 \end{bmatrix} \mathbf{u}_n^c \ . \qquad (5)$$

As discussed, the motion of the system is split into pure rotation and pure translation. The required orientations of the robots $\theta_n^{r_i}$ and robot controls $\mathbf{u}_n^{r_i}$ change depending on the centroid's motion. The sign in Eq. (4) depends on the direction of the centroid's rotation. The first and second columns of the matrix in Eq. (5) correspond to the centroid's translation and rotation, respectively. Depending on the specific formation and system dynamics of a system, Eq. (5) can be modified and extended. The conversion between the centroid velocity $\mathbf{v}_n^c$ and centroid control $\mathbf{u}_n^c$ is trivial and omitted for conciseness. These reductions allow for efficient optimization while improving scalability dramatically. It also reduces the inter-robot communication while still able to achieve impressive performance, as shown in experiments.

B. Factor Graphs

We use factor graphs to perform efficient and unified optimization of the joint control and planning problem. Factor graphs are a bipartite graphical structure that represents a factorized function, such that variable relationships are encapsulated by edges and factors. We represent our path optimization problem as a MAP inference problem, and use GTSAM [48], to perform inference over the factor graph.

We can represent the centroid poses, $\mathbf{x}^c$, and the centroid velocity $\mathbf{u}^c$ as variables in the factor graph. We develop relationships between these variables and use sparse optimizers to solve for our control efforts. We found that Levenberg–Marquardt (LM) is a solver well-suited for most of our use cases, as it provides optimization results of decent accuracy and great robustness. LM solver provides fast and local convergence guarantees [49], which are important to real-time systems. As shown in Fig. 3, factors used are: f_n^x unary pose factor, penalizes the system for deviating away from the provided initial path; f_n^u unary speed factor, constrains the speed at which the system moves along the planned path to the motor speed limit; f_n^m ternary motion factor, forces the algorithm to adhere to the established motion model; and $f_n^{obs_j}$ obstacle avoidance factor, penalizes the system for a path nearing obstacles. The overall cost function is

$$J(\mathbf{x}, \mathbf{u}) = \sum_{n=k}^{N} e_n^x + \sum_{n=k}^{N-1} (e_n^u + e_n^m) + \sum_{j=1}^{J} \sum_{n=k+1}^{N} e_n^{obs_j} \ , \quad (6)$$

$$e_n^x(\mathbf{x}_n^c) = ||\mathbf{x}_n^c - \mathbf{x}_n^{c,ref}||_{\Omega_n^x}^2 \ , \quad (7)$$

$$e_n^u(\mathbf{u}_n^c) = ||\mathbf{u}_n^c - \mathbf{u}_n^{c,ref}||_{\Omega_n^u}^2 \ , \quad (8)$$

$$e_n^m(\mathbf{x}_n^c, \mathbf{u}_n^c, \mathbf{x}_n^c) = ||\mathbf{x}_{n+1}^c(\mathbf{x}_n^c, \mathbf{u}_n^c)^* - \mathbf{x}_{n+1}^c||_{\Omega_n^m}^2 \ , \quad (9)$$

$$e_n^{obs_j}(\mathbf{x}_n^c) = \begin{cases} ||1 - \frac{d_j}{R}||_{\Omega_n^{obs}}^2, & d_j < R \\ 0, & d_j \geq R \end{cases} \ . \quad (10)$$

$||.||$ is the Mahalanobis distance, with Ω terms representing the information matrices of the variables. e_n^x and e_n^u are the costs corresponding to factors f_n^x and f_n^u, respectively. e_n^m is the cost corresponding to factors f_n^m, where $\mathbf{x}_{n+1}^c(\mathbf{x}_n^c, \mathbf{u}_n^c)^*$ represents the estimated pose $\mathbf{x}_{n+1}^{c}$* to which the centroid moves over the next time step provided the optimized pose $\mathbf{x}_n^c$ and control $\mathbf{u}_n^c$. e_n^{obs} is the defined cost corresponding to factors f_n^{obs}, where d_j represents the centroid's distance from obstacle j and R represents the radius of a predefined safety bubble that encapsulates the entire multi-robot system [30]. Factors can be added, removed, or modified depending on the task requirements. Subtractions are overloaded by appropriate manifold operations handled in GTSAM.

With proper selection of the covariances, the equivalent effects of regulation matrix adjustments can be achieved and terms that are commonly considered constraints (motion model and obstacle) can be well approximated with soft constraints as have been shown in similar factor graph implementations in [30] and [21]. It was also proven in [50] that covariances of 0 yield deterministic solutions that satisfy the boundaries confined by the constraints. There are other alternatives utilizing barrier functions [51], Lagrange multipliers, and dual


Fig. 3: Multi-robot control and planning factor graph with centroid pose variables and centroid control variables. Factors and their associated costs are defined in the text. Since estimation is not the emphasis of this work, only the variables at the current state and future states are used in the optimization problem for each iteration. A finite horizon is used in our case, extending to the terminal state to ensure long-term stability.

graph optimization [52], but we find the formulation presented above sufficient for our application.

To analyze the complexity, we perform a worst-case complexity analysis on our approach and MPC. For our approach, we create a factor graph with $O(2n + 4n)$ complexity, 2 variables, and 4 factors are added at each step, n. The LM solver with multifrontal QR decomposition can be approximated $\tilde{O}(\epsilon^{-2})$ [53] where $\tilde{O}(.)$ indicates the presence of logarithmic factors in ϵ. Post-optimization extractions take $O(I)$, as such, $\tilde{O}(n\epsilon^{-2}(n^2 + m) + nI)$ where m are solver parameters. We can develop the MPC formulation in a similar manner, with the complexity of Augmented Lagrangian solver taken as $O(|\log(\epsilon)|^2 \epsilon^{-2})$ [54], the complexity becomes $O(nh(|\log(\epsilon)|^2 \epsilon^{-2}) + nI)$, where h are the horizon steps. Both methods demonstrate polynomial scaling with the number of steps and an inverse-square relationship with error tolerance. However, our emphasis on scalability concerns the number of robots and environmental complexity rather than just task steps, N. In these aspects, our approach exhibits a linear behavior, and yields shorter optimization times compared to MPC, effectively enhancing performance for larger numbers of robots and more complex environments.

IV. EXPERIMENTS

Experiments are performed in three environments: i) a pure-simulated environment, ii) a physics-based Gazebo simulation [55], and iii) real-world hardware. The performances are evaluated by five metrics: 1) average deviation, 2) inter-robot errors, 3) optimization time, 4) path length, and 5) distance to goal. The average deviation indicates the average distance between each state and its corresponding state on the *initial path* estimate. The *initial path* is a straight-line path from the start to the goal point. It will be optimized eventually. We assume that initial robot formation is performed beforehand, and as such, the inter-robot error describes the amount by which distances between robots change when no payloads are present. This measures the impact of Eq. (5) on the system. The optimization time measures the time taken to optimize the paths and obtain individual robot control efforts. All

TABLE I: Table outlining the benefits of using penalty (MPC-P) vs equalities (MPC-C) for the MPC approach. We selected the option that provided the best advantage to MPC. We also compare gradient (G) to gradient-free approaches (GF).

	MPC-C		MPC-P	
	G	GF	G	GF
Time (ms)	34.93	89.28	23.95	55.584
Deviation (m)	0.02	0.02	0.08	0.10
Distance to Goal (m)	0.003	0.028	0.006	0.062
Path Length (m)	9.04	8.99	9.37	9.15
Obstacle Free	✓	✗	✓	✗

experiments were run on an Intel i7-10700 at 4.8 GHz, with 16 GB DDR4 memory. The path length is a metric for assessing the system's ability to reach the end state following the shortest path, even when perturbed or blocked by an obstacle. The distance to goal metric evaluates the system's ability to reach the goal and ensures short paths do not get over-prioritized.

A. Experiment 1: Simulation

Simulations in Experiment 1 are designed to test the scalability by increasing the number of robots, from 4 to 128, and the complexity of the environment by introducing more obstacles. It is implemented in C++, with metrics including optimization time, path deviation, and path length.

Here the baseline is model predictive control (MPC), based on NLopt [56]. We implement two MPC algorithms; gradient and gradient-free approaches: augmented Lagrangian AUGLAG [7] and constrained optimization by linear approximation COBYLA [8]. The MPC was designed to handle system requirements as either constraints or penalties (penalties are implemented similarly to our approach), where the former provides more accurate results and the latter optimizes faster. We define these as MPC-C and MPC-P, respectively. As shown in Table I, both approaches are used in our comparative analysis. The best parameters for all algorithms were chosen by performing an exhaustive search over all the parameters within 2000 iterations (see supplementary), such that the distance to the goal is below $0.06m$ and the optimization time does not exceed $80ms$. In total, 20580 parameter sets were examined. See Table II for the best parameters.

Robots traversed from (-2, 0) to (7, 0) with obstacles at (1, 0.4), (2, -0.4), (3, -0.6), (4, 0.7) (5, 0.4), all in meters. We measured optimization time, average deviation from the initial path, and the path length. Results are plotted in Fig. 4.


Fig. 4: Results of Experiment 1 showing the scalability with respect to the number of (left) robots, and (right) obstacles.

The MPC results in the figure are obtained from MPC-C and MPC-P with gradient-based solvers, as they have been proven to be faster and more reliable (see Table I).

As shown in Fig. 4, under the reduced formulation of the problem, the factor graph representation consistently outperforms MPC-C and MPC-P in all experiments. The optimization time of the three algorithms all scaled linearly in the robot scalability experiment, Fig. 4-(left). Such increases in times are associated with some overhead computing individual robot controls that increase with the number of robots which are minimal and can be further optimized.

It can also be seen from Fig. 4-(right) that the factor graph representation has allowed the optimization to scale better as more obstacles are added to the environment. We notice the optimization time being 38.4%, 37.9%, 52.1% and 76.1% shorter than those of MPC-C, and 17.1%, 12.2%, 23.7% and 56.1% shorter than those of MPC-P in this obstacle scalability experiment. Our approach consistently reaches the goal, while MPC-C's distances to the goal are $0.003m$, $0.001m$, $0.002m$, and $0.002m$, and MPC-P's $0.002m$, $0.006m$, $0.006m$, and $0.01m$. One downside to our approach is that it yields path lengths longer than MPC-C's by 2.99%, 3.72%, and 7.16%, and 6.56% and MPC-P's by 1.73%, 0.81%, 3.38%, and 3.06%. Increasing obstacles from 1 to 2, the optimization times for our approach, MPC-C, and MPC-P increased by 10.1%, 9.2%, and 3.95%. Increasing from 2 to 7, the optimization times increased by another 5.84%, 174.4%, and 111.5%

B. Experiment 2: Gazebo Simulation

In addition to MPC, we compare with another baseline, hierarchical quadratic programming (HQP) [5]. Since the code for HPQ is not open-source, we used the same Gazebo [55] world with identical start/goal positions and geometric parameters described in [5] to run our algorithm. We compare inter-robot errors, as deviation and temporal results were not provided in [5]. As shown in Table III, our approach outperforms

TABLE II: The best parameters for the algorithms by exhaustive search. The parameters serve as weights/covariances onto the associated errors in the cost function.

Parameters	Ours (Covariance)	MPC-C (Weight)	MPC-P (Weight)
State	$(1m^2, 1m^2, 0.2rad^2)\times 10^{-1}$	$1m^{-2}$	$1m^{-2}$
Terminal	$(1m^2, 1m^2, 0.2rad^2)\times 10^{-1}$	$1\times 10^3 m^{-2}$	$1m^{-2}$
Control	$(1m^2 s^{-2}, 1m^2 s^{-2})\times 10^{-1}$	$1\times 10^{-3} m^{-2} s^2$	$1m^{-2} s^2$
Motion Model	$(1m^2, 1m^2, 0.2rad^2)\times 10^{-4}$	–	$1\times 10^{-1} m^{-2}$
Obstacle	0.01	–	1
Horizon	–	2 steps	2 steps
Relative Tolerance	1×10^{-2}	1×10^{-4}	1×10^{-2}
Absolute Tolerance	1×10^{-2}	1×10^{-4}	1×10^{-2}
Error Tolerance	1×10^{-2}	1×10^{-4}	1×10^{-2}

TABLE III: Simulation Results: Payload navigating from (0,0) to (5,3) in Gazebo environment. The length of the initial path estimate is 5.831m. MPC-C (constraint), MPC-P (penalty), and hierarchal quadratic programming HQP [5] are shown.

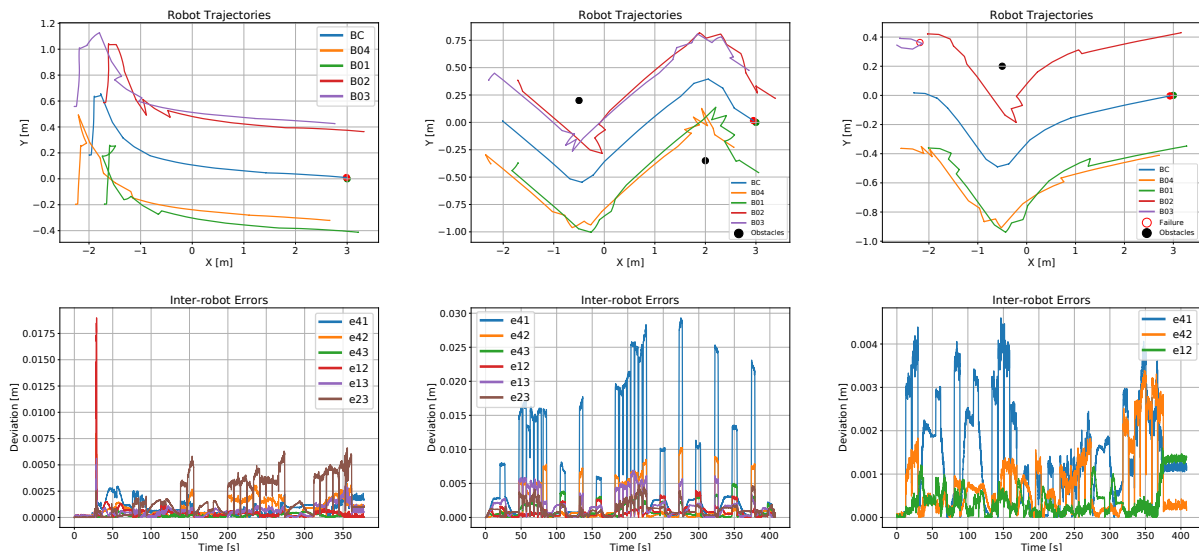|	Ours	MPC-C	MPC-P	HQP [5]
Avg. Deviation (m)	**0.015**	0.017	0.019	-
Max Inter-Robot Error (m)	**0.016**	0.05	0.05	0.017
Path Length (m)	**5.824**	5.825	5.825	-
Dist. to Goal (m)	**0.017**	0.026	0.023	-

Fig. 5: Hardware trial results. A payload was pushed by four Turtlebot3 Waffle Pi robots from random start points near the origin (0, 0) to the same goal (3, 0). The three columns correspond to the three types of tasks: (left column) reaching the goal while undergoing large disturbances, (middle column) obstacle avoidance, and (right column) robot failure during obstacle avoidance. The system is consistently able to recover and reach the goal within 1cm accuracy.

HQP in terms of inter-robot errors, and MPC in average path deviation, inter-robot error, and distance to goal.

C. Experiment 3: Real-world Robots

As shown in Fig. 1, four TurtleBot3 Waffle Pi robots are deployed with an object lifted on top of them. These robots are connected through revolute joints to the object, as they allow for fixed geometry of the system while also maintaining the system's ability to fully rotate. ROS2 is used as the communication middleware. We obtain robot poses from a Vicon motion capture system with 1mm accuracy. We present three scenarios to evaluate the performance and demonstrate the robustness of the framework: i) disturbances, ii) obstacle avoidance, and iii) robot failure. The results are presented in Fig. 5. The rows show robot paths and the inter-robot errors. Fig. 5-(left column) shows the case where the system was disturbed by 40cm, at time 30 sec, while executing a path and was able to recover from the disturbance within a few steps to reach the goal with sub-1 cm accuracy. Fig. 5-(middle column) shows the system tasked with obstacle avoidance while on a mission, and successfully avoided the obstacles reaching the goal point within 1cm accuracy. Fig. 5-(right column) shows a robot failure. One robot was disabled and removed from the trial. The system was able to compensate and reach the goal within 1cm accuracy.

V. CONCLUSIONS

In this paper, we proposed a novel technique for multi-robot joint planning and control, leveraging the scalability and robustness of factor graphs and the inherent forward-looking nature of the planning and control problems. We make use of a single optimization to reduce optimization times, which is further aided by reducing the dimensionality

to yield a highly scalable approach. We compare our approach against state-of-the-art and consistently outperformed previous approaches. Furthermore, our approach is highly extensible. Adding new constraints on the system is a matter of simply adding factors to the graph, and dynamics can be modified through simple modification of the motion model while the optimization remains the same. Comparative results show an improvement in optimization time, observing a reduction of 50.1% in scalability experiments. Our algorithm consistently reduced the average path deviation and optimization time and produced a shorter distance to the goal. Similarly, our hardware experiments showed the robustness of our approach, with demonstrations of disturbances, obstacle avoidance, and robot failures. Scalability experiments showed that the computation time does not significantly increase with increased environment complexity. This is due to both the formulation and the use of factor graphs; as MPC approaches made use of similar formulation and observed similar scaling, while the factor graph allowed us to outperform MPC.

In future works, we will integrate Gaussian processes into the control problem to allow for smoother paths and a continuous-time representation. Another direction is extending the work to motions with higher degrees of freedom.

REFERENCES

[1] L. E. Parker, *Multiple Mobile Robot Teams, Path Planning and Motion Coordination*, pp. 5783–5800. Springer, 2009.

[2] M. Mukadam, J. Dong, F. Dellaert, and B. Boots, "Steap: simultaneous trajectory estimation and planning," *Autonomous Robots*, vol. 43, no. 2, pp. 415–434, 2018.

[3] T. Halsted, O. Shorinwa, J. Yu, and M. Schwager, "A survey of distributed optimization methods for multi-robot systems," *ArXiv*, vol. abs/2103.12840, 2021.

[4] G. Habibi, Z. Kingston, W. Xie, M. Jellins, and J. McLurkin, "Distributed centroid estimation and motion controllers for collective transport by multi-robot systems," in *ICRA*, pp. 1282–1288, 2015.

[5] D. Koung, O. Kermorgant, I. Fantoni, and L. Belouaer, "Cooperative multi-robot object transportation system based on hierarchical quadratic programming," *IEEE RA-L*, vol. 6, no. 4, pp. 6466–6472, 2021.

[6] F. Dellaert and M. Kaess, *Factor Graphs for Robot Perception*. Foundations and Trends in Robotics Series, Now Publishers, 2017.

[7] A. R. Conn, N. I. M. Gould, and P. Toint, "A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds," *SIAM Journal on Numerical Analysis*, vol. 28, pp. 545–572, 1991.

[8] M. J. D. Powell, *A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation*, pp. 51–67. Dordrecht: Springer Netherlands, 1994.

[9] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE SMC*, vol. 4, no. 2, pp. 100–107, 1968.

[10] H. M. Choset, S. Hutchinson, K. M. Lynch, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of Robot Motion: theory, algorithms, and implementation*. MIT press, 2005.

[11] S. Lavalle and J. Kuffner, "Rapidly-exploring random trees: Progress and prospects," *Algorithmic and Computational Robotics*, 2000.

[12] L. E. Kavralu, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE T-RO*, 1996.

[13] M. Zucker, N. Ratliff, A. D. Dragan, M. Pivtoraiko, M. Klingensmith, C. M. Dellin, J. A. Bagnell, and S. S. Srinivasa, "Chomp: Covariant hamiltonian optimization for motion planning," *IJRR*, vol. 32, no. 9-10, pp. 1164–1193, 2013.

[14] A. Byravan, B. Boots, S. Srinivasa, and D. Fox, "Space-time functional gradient optimization for motion planning," in *ICRA*, pp. 6499 – 6506, 05 2014.

[15] K. He, E. Martin, and M. Zucker, "Multigrid chomp with local smoothing," in *2013 13th IEEE-RAS (Humanoids)*, pp. 315–322, 2013.

[16] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "Stomp: Stochastic trajectory optimization for motion planning," in *2011 IEEE ICRA*, pp. 4569–4574, 2011.

[17] M. Toussaint, "Robot trajectory optimization using approximate inference," in *ICML*, p. 1049–1056, 2009.

[18] J. Dong, M. Mukadam, F. Dellaert, and B. Boots, "Motion planning as probabilistic inference using gaussian processes and factor graphs.," in *RSS*, vol. 12, 2016.

[19] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *IJRR*, vol. 33, no. 9, pp. 1251–1270, 2014.

[20] M. Mukadam, X. Yan, and B. Boots, "Gaussian process motion planning," in *2016 IEEE (ICRA)*, pp. 9–15, 2016.

[21] B. Bazzana, T. Guadagnino, and G. Grisetti, "Handling Constrained Optimization in Factor Graphs for Autonomous Navigation," 2022.

[22] M. Xie, A. Escontrela, and F. Dellaert, "A factor-graph approach for optimization problems with dynamics constraints," *arXiv preprint arXiv:2011.06194*, 2020.

[23] C. Park, J. Pan, and D. Manocha, "ITOMP: Incremental trajectory optimization for real-time replanning," in *ICAPS*, 2012.

[24] A. Patwardhan and A. J. Davison, "A distributed multi-robot framework for exploration, information acquisition and consensus," *arXiv preprint arXiv:2310.01930*, 2023.

[25] R. Murai, J. Ortiz, S. Saeedi, P. H. J. Kelly, and A. J. Davison, "A Robot Web for Distributed Many-Device Localisation," *IEEE T-RO*, vol. 40, pp. 121–138, 2024.

[26] I. Spasojevic, X. Liu, A. Ribeiro, G. J. Pappas, and V. Kumar, "Active Collaborative Localization in Heterogeneous Robot Teams," *arXiv preprint arXiv:2305.18193*, 2023.

[27] L. Guerrero-Bonilla, A. Prorok, and V. Kumar, "Formations for Resilient Robot Teams," *IEEE RA-L*, vol. 2, no. 2, pp. 841–848, 2017.

[28] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized multi-robot path planning," in *IROS*, pp. 11785–11792, 2020.

[29] M. Cao, K. Cao, S. Yuan, K. Liu, Y. L. Wong, and L. Xie, "Path Planning for Multiple Tethered Robots Using Topological Braids," *RSS*, 2023.

[30] A. Patwardhan, R. Murai, and A. J. Davison, "Distributing collaborative multi-robot planning with gaussian belief propagation," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 552–559, 2023.

[31] A. Pierson, A. Ataei, I. C. Paschalidis, and M. Schwager, "Cooperative multi-quadrotor pursuit of an evader in an environment with no-fly zones," in *ICRA*, pp. 320–326, May 2016.

[32] D.-N. Ta, M. Kobilarov, and F. Dellaert, "A factor graph approach to estimation and model predictive control on unmanned aerial vehicles," in *2014 ICUAS*, pp. 181–188, 2014.

[33] A. D. Saravanos, Y. Li, and E. A. Theodorou, "Distributed Hierarchical Distribution Control for Very-Large-Scale Clustered Multi-Agent Systems," *RSS*, 2023.

[34] L. Yan, T. Stouraitis, and S. Vijayakumar, "Decentralized Ability-Aware Adaptive Control for Multi-Robot Collaborative Manipulation," *IEEE RAL*, vol. 6, no. 2, pp. 2311–2318, 2021.

[35] Y. He, M. Wu, and S. Liu, "An optimisation-based distributed cooperative control for multi-robot manipulation with obstacle avoidance," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 9859–9864, 2020.

[36] Y.-H. Liu, S. Arimoto, and T. Ogasawara, "Decentralized cooperation control: Non-communication object handling," in *ICRA*, vol. 3, pp. 2414–2419 vol.3, 1996.

[37] D. Rus, B. Donald, and J. Jennings, "Moving furniture with teams of autonomous robots," in *IROS*, vol. 1, pp. 235–242 vol.1, 1995.

[38] M. Lujak, "A Distributed Coordination Model for Multi-Robot Box Pushing," *IFAC Proceedings Volumes*, vol. 43, no. 4, pp. 120–125, 2010.

[39] J. Alonso-Mora, R. Knepper, R. Siegwart, and D. Rus, "Local motion planning for collaborative multi-robot manipulation of deformable objects," in *ICRA*, pp. 5495–5502, 2015.

[40] Z. Wang and M. Schwager, "Multi-robot manipulation with no communication using only local measurements," in *CDC*, pp. 380–385, 2015.

[41] Z. Wang and M. Schwager, "Kinematic multi-robot manipulation with no communication using force feedback," in *ICRA*, pp. 427–432, 2016.

[42] Z. Wang and M. Schwager, "Force-amplifying n-robot transport system (force-ants) for cooperative planar manipulation without communication," *IJRR*, vol. 35, no. 13, pp. 1564–1586, 2016.

[43] P. Culbertson and M. Schwager, "Decentralized Adaptive Control for Collaborative Manipulation," in *ICRA*, pp. 278–285, 2018.

[44] R. T. Fawcett, L. Amanzadeh, J. Kim, A. D. Ames, and K. A. Hamed, "Distributed data-driven predictive control for multi-agent collaborative legged locomotion," in *ICRA*, pp. 9924–9930, 2023.

[45] R. T. Fawcett, K. Afsari, A. D. Ames, and K. A. Hamed, "Toward a data-driven template model for quadrupedal locomotion," *IEEE RA-L*, vol. 7, no. 3, pp. 7636–7643, 2022.

[46] O. Shorinwa and M. Schwager, "Scalable collaborative manipulation with distributed trajectory planning," in *IROS*, pp. 9108–9115, 2020.

[47] O. Shorinwa and M. Schwager, "Distributed contact-implicit trajectory optimization for collaborative manipulation," in *MRS*, pp. 56–65, 2021.

[48] Frank Dellaert and GTSAM Contributors, "GTSAM." https://github. com/borglab/gtsam, May 2022.

[49] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters," *SIAM*, vol. 11, no. 2, pp. 431–441, 1963.

[50] S. Yang, G. Chen, Y. Zhang, H. Choset, and F. Dellaert, "Equality constrained linear optimal control with factor graphs," in *ICRA*, IEEE, May 2021.

[51] R. Grandia, F. Farshidian, R. Ranftl, and M. Hutter, "Feedback MPC for torque-controlled legged robots," in *IROS*, pp. 4730–4737, 2019.

[52] M. Xie, A. Escontrela, and F. Dellaert, "A factor-graph approach for optimization problems with dynamics constraints," *ArXiv*, vol. abs/2011.06194, 2020.

[53] E. H. Bergou, Y. Diouane, and V. Kungurtsev, "Convergence and Complexity Analysis of a Levenberg–Marquardt Algorithm for Inverse Problems," *Journal of Optimization Theory and Applications*, vol. 185, pp. 927–944, June 2020.

[54] G. N. Grapiglia and Y.-x. Yuan, "On the complexity of an augmented Lagrangian method for nonconvex optimization," *IMA Journal of Numerical Analysis*, vol. 41, pp. 1546–1568, 07 2020.

[55] Open Robotics, "Gazebo: Simulate before you build." https://gazebosim. org/home, 2023.

[56] S. G. Johnson, "The NLopt nonlinear-optimization package." http:// github.com/stevengj/nlopt, 2007.