

# Control-Tree Optimization: an approach to MPC under discrete Partial Observability

Camille Piquepal<sup>1</sup> and Marc Toussaint<sup>2</sup>

**Abstract**—This paper presents a new approach to Model Predictive Control for environments where essential, discrete variables are partially observed. Under this assumption, the belief state is a probability distribution over a finite number of states. We optimize a *control-tree* where each branch assumes a given state-hypothesis. The control-tree optimization uses the probabilistic belief state information. This leads to policies more optimized with respect to likely states than unlikely ones, while still guaranteeing robust constraint satisfaction at all times. We apply the method to both linear and non-linear MPC with constraints. The optimization of the *control-tree* is decomposed into optimization subproblems that are solved in parallel leading to good scalability for high number of state-hypotheses. We demonstrate the real-time feasibility of the algorithm on two examples and show the benefits compared to a classical MPC scheme optimizing w.r.t. one single hypothesis.

## I. INTRODUCTION

In the field of receding horizon motion planning, uncertainty about the robot environment is often neglected or integrated in the environment representation through heuristics (e.g. by adding safety distances, or potential fields around perceived objects). At each planning cycle one single environment representation is used for planning the next motion. Reactivity and robustness are achieved by re-planning.

In contrast, we aim here to tackle problems where the uncertainty is about discrete and critical aspects of the environment, such that the agent has to anticipate for multiple hypothetical scenarios within the planning horizon.

The proposed approach is inspired from [1] where an integrated Task and Motion planner is introduced for partial observable environments. [1] focuses mainly on off-line planning for object manipulation and optimizes trajectory-trees that react to observation. This paper transfers this concept to Model Predictive Control and real-time trajectory optimization. Accordingly, the main contributions of this paper are:

- An MPC formulation for optimizing control policies (*control-trees*) under discrete partially observability. It leverages the probabilistic information of the belief state to reduce conservativeness.
- An algorithm *Distributed Augmented Lagrangian* combining the Augmented Lagrangian and ADMM method. It allows fast and distributed optimization of the control-tree, is applicable to both linear and non-linear MPC, and scales linearly w.r.t. the number of state hypotheses.

- The demonstration of the benefits and real-time feasibility of the method on robotics problems.

## II. RELATED WORK

Stochastic and Robust variants of Model Predictive Control aim at optimizing performance and constraint satisfaction under disturbances (e.g. un-modeled system dynamics, external disturbances, noise). Good surveys on Stochastic and Robust Optimal can be found in [2] and [3]. Robust MPC seeks constraints satisfaction at all times by tackling the worst cases of the uncertainty. In contrast, Stochastic MPC uses the information how the uncertainty is distributed and interprets constraints probabilistically: a constraint is satisfied if the probability of a violation is kept below some threshold. Like Robust MPC, the presented approach seeks constraint satisfaction even in the worst case. As per Stochastic MPC, it leverages the stochastic information to reduce the conservativeness.

Several approaches have been developed that consider multiple scenarios within the planning horizon.

Min-Max MPC [4][5] can be considered as one of these, which optimizes a control sequence w.r.t. the worst case. Although this guarantees that constraints are met at all times, it can be over-conservative.

Approaches optimizing control-trees have already been used in MPC, mainly for optimal control of chemical processes. In [6][7][8] Lucia and al. introduced a tree-based approach called *multi-stage* MPC. This has been developed further, e.g. Kouzoupis and al. in [9], Leidreiter and al. [10].

To our knowledge [11] and [12] are the only published experiments applying *multi-stage* MPC to mobile robotic use-cases. In all those previous applications of multi-stage MPC, scenarios are used to model uncertainty over continuous variables (by enumerating possible realizations). Its application as a treatment of partial observability, with intergration of the belief state is novel. Moreover, unlike [12], the approach is not restricted to linear MPC, and compared to [11], we add an optimization algorithm allowing distributed computation to scale to high number of hypotheses.

## III. PROBLEM FORMULATION

### A. Overview

The presented approach optimizes control policies in a context of mixed-observability. The usual continuous state variables (e.g. position, velocity) are assumed fully observed. On the other hand, discrete, essential parts of the environment are partially observed, this is modelled by a discrete

<sup>1</sup>Machine Learning & Robotic Lab, University of Stuttgart, Germany  
camille.piquepal@ipvs.uni-stuttgart.de

<sup>2</sup>Learning and Intelligent Systems Lab, TU Berlin, Germany  
toussaint@tu-berlin.de

state. In the example of Figure 1), this models the pedestrian intention.

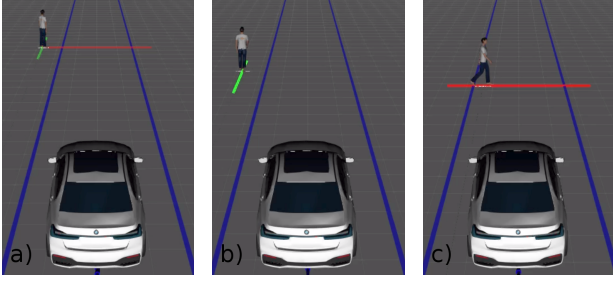


Fig. 1: In a) The pedestrian is standing, his intention is uncertain. He may walk along the street b) or cross c). It is modeled as a partially observable discrete state.

We assume that the agent is not oblivious to the likelihood of each hypothesis: it maintains a belief state, i.e. a probability distribution over the multiple hypotheses. The belief state is used to optimize the policy more w.r.t. likely states.

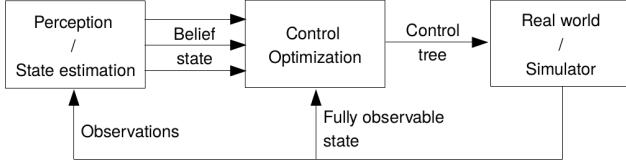


Fig. 2: Control loop: control-trees are optimized w.r.t. a belief state (probability distribution over multiple hypotheses).

### B. State representation

We consider a compound state representation where a state is composed of 2 parts:

- $x \in \mathbb{R}^n$ : system continuous state
- $s \in S$ : discrete state from a finite state space  $S$

The continuous state  $x$  corresponds to the classical notion of state in the MPC literature, and is assumed fully observable.

The discrete state  $s$  represents discrete aspects of environment and is only partially observed.

### C. Control-tree

The algorithm consists in optimizing a control-tree. The control-tree has a first common trunk (see orange part in Figure 3). This part is executed by the controller until the next planning cycle happens. We call *control horizon* the time interval corresponding to this common trunk. After the control horizon, the control-tree branches, and each branch is optimized assuming a certain state  $s$ . This is akin to the QMDP method (see [13]) for solving POMDPs, but here in continuous domain.

### D. Optimization problem formulation

We note  $L$  the number of steps in the control horizon, and  $T$  the total number of steps on each branch ( $L \ll T$ ).

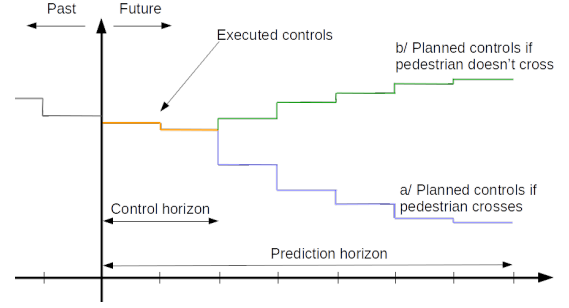


Fig. 3: Schematic control-tree applied to the pedestrian example: After the control horizon, the control-tree evolves in two different branches corresponding to the two different hypothetical states (pedestrian crossing or not).

Let  $u_s$  and  $x_s$ ,  $s \in S$  be the control and configuration sequences on the branch corresponding to the state  $s$ . Let  $\tilde{u}$  be the controls in the control horizon. Let  $p(s)$  be the probability of state  $s$ .

We formulate the optimal control problem as follows:

$$\min_{\tilde{u}, u_s, x_s} \sum_{s \in S} p(s) \sum_{t=0}^{T-1} c_s(x_s(t), u_s(t)), \quad (1a)$$

s.t.

$$g_s(x_s(t), u_s(t)) \leq 0, \quad \forall s, \forall t, \quad (1b)$$

$$x_s(t+1) = f(x_s(t), u_s(t)), \quad \forall s, \forall t, \quad (1c)$$

$$u_s(t) = \tilde{u}(t), \quad \forall s, \forall t < L, \quad (1d)$$

where  $c_s$  and  $g_s$  are the cost and inequality constraints applying in state  $s$  respectively. The function  $f$  models the system dynamic.

The optimization variables  $u_s$  and  $x_s$ ,  $s \in S$  are not fully independent from each other. Indeed, in the control horizon, they all correspond to the common trunk and must therefore be equal. This is captured by the equation (1d) and is usually called "non-anticipativity" constraint in multi-stage MPC [6]. In the following we call  $\tilde{u}$  the consensus variable.

The common trunk of the tree is constrained by the active constraints of all states (see (1b)), regardless of the state likelihood. This is for guaranteeing the robustness of the constraint satisfaction.

On the other hand, the minimized costs (1a) are weighted by the belief state for optimizing more w.r.t. likely states than unlikely ones.

### E. Transcription to generic solver format

Here we rewrite the optimization problem in a generic format that is the input to our solver. In Equation (1) both the controls  $u_s$  and states  $x_s$  are optimization variables. In many cases it is possible to eliminate either the controls or the configurations and obtain a more compact formulation:

- Optimization in control space. This is achieved by eliminating the  $x_s$  and is well described in the MPC literature, in particular for linear MPC (e.g. [14]). We apply this scheme in the first experiment in section V-A.

- Optimization in configuration space. This typically requires adding additional constraints to ensure the existence of controls implementing the configuration transitions e.g. for non-holonomic robots. The second experiment in section V-B follows this approach.

The problem can be rewritten in a generic format:

$$\min_{\tilde{z}, z_s} \sum_{s \in S} p(s) c_s(z_s), \quad (2a)$$

s.t.

$$g_s(z_s) \leq 0, \quad \forall s, \quad (2b)$$

$$h_s(z_s) = 0, \quad \forall s, \quad (2c)$$

$$z_s(t) = \tilde{z}(t), \quad \forall s, \forall t < L, \quad (2d)$$

where  $\tilde{z}$  and  $z_s, s \in S$  are the optimization variables, which can be in control space, configuration space, or a combination of both. We note  $d$  the dimensionality of the optimized parameters at each time step. For each  $s \in S$ , the function  $c_s : \mathbb{R}^{T \times d} \rightarrow \mathbb{R}$  is a scalar objective function,  $g_s : \mathbb{R}^{T \times d} \rightarrow \mathbb{R}^{d_{g_s}}$  and  $h_s : \mathbb{R}^{T \times d} \rightarrow \mathbb{R}^{d_{h_s}}$  define  $d_{g_s}$  inequality and  $d_{h_s}$  equality constraints respectively. (2d) is the non-anticipativity constraint and  $\tilde{z}$  is the consensus variable. We generally assume the functions  $c_s$ ,  $g_s$ , and  $h_s$  to be smooth, but not necessarily convex or unimodal.

#### IV. SOLVER

The global optimization problem (2) can be decomposed into  $N$  loosely coupled optimization problems with  $N = |S|$ . Indeed the  $N$  tree branches are nearly-independent from each other (except for the non-anticipativity constraint (2d)). The core idea of the proposed algorithm is to take advantage of this decomposition. It performs multiple iterations consisting each of 2 phases:

- A distributed phase where a relaxed version of each subproblem is optimized.
- A centralized phase at which the consensus variable  $\tilde{z}$  is updated.

##### A. Distributed Augmented Lagrangian

The optimization problem of each branch is nearly a standard constrained optimization problem. Only the non-anticipativity constraint makes it peculiar, because the consensus variable  $\tilde{z}$  varies over time and depends on the results of the other subproblems.

To solve this, we form an unconstrained objective function which combines the augmentations of both the Augmented Lagrangian method (AuLa) and the Alternating Direction Method of Multipliers (ADMM). We call it the *Distributed Augmented Lagrangian*:

$$L_s(z, \tilde{z}, \lambda_s, \kappa_s, \eta_s) = p(s) c_s(z) \quad (3a)$$

$$+ \lambda_s^T g_s(z) + \mu \| [g_s(z) > 0] \cdot g_s(z) \|^2 \quad (3b)$$

$$+ \kappa_s^T h_s(z) + \nu \| h_s(z) \|^2 \quad (3c)$$

$$+ \eta_s^T (\Delta z) + \frac{\rho}{2} \|\Delta z\|^2, \quad (3d)$$

where  $\Delta z \in \mathbb{R}^{L \times d}$ , with  $\Delta z(t) = z(t) - \tilde{z}(t), \forall t < L$  is the difference between  $z$  and the consensus  $\tilde{z}$  in the control horizon.

(3b) and (3c) are the terms of the Augmented Lagrangian for handling the inequality and equality constraints intrinsic to each subproblem (see [15], [16]).

(3d) are the similar ADMM terms for solving the coupling between the subproblems (see [17]).

$\lambda_s \in \mathbb{R}^{d_{g_s}}, \kappa_s \in \mathbb{R}^{d_{h_s}}, \eta_s \in \mathbb{R}^{L \times d}$  are the dual variables, and  $\mu, \nu, \rho$  are fixed real constants.

##### B. Optimization procedure

The optimization algorithm is the following:

$$z_s^{k+1} := \underset{z_s}{\operatorname{argmin}} L_s(z_s^k, \tilde{z}^k, \lambda_s^k, \kappa_s^k, \eta_s^k) \quad (4a)$$

$$\lambda_s^{k+1} := \max(0, \lambda_s^k + 2\mu g_s(z_s^{k+1})) \quad (4b)$$

$$\kappa_s^{k+1} := \kappa_s^k + 2\nu h_s(z_s^{k+1}) \quad (4c)$$

$$\tilde{z}^{k+1} := \frac{1}{N} \sum_{s \in S} z_s^{k+1} \quad (4d)$$

$$\eta_s^{k+1} := \eta_s^k + \rho \Delta z_s^{k+1}. \quad (4e)$$

1) *Initialization*: All dual variables  $\lambda_s, \kappa_s, \eta_s$  are initially set to 0.0. The optimization variables  $z_s$  as well as the initial reference  $\tilde{z}$  can be initialized randomly, or by any better heuristic to speed-up.

2) *Unconstrained minimization*: The step (4a) is the minimization of  $N$  unconstrained optimization problems and is the part which is computationally expensive. We optimize them with a newton procedure. As Figure 4 shows, these optimizations can all be performed in parallel. This is where the algorithm takes full advantage of the decomposition.

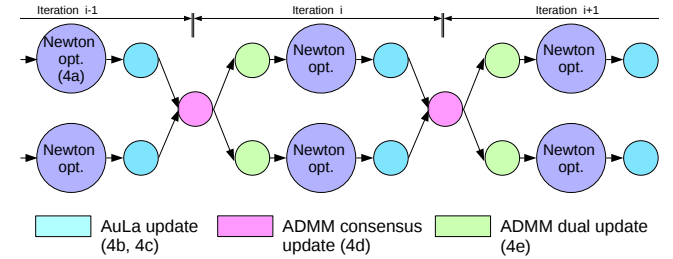


Fig. 4: Execution flow for  $N = 2$ . The costly steps (Newton minimizations) are parallelized.

3) *Augmented Lagrangian dual variables update*: Equations (4b) and (4c) update the dual variables corresponding to the inequality and equality constraints respectively.

4) *ADMM dual and consensus variables update*: Line (4d) updates the consensus variable. It can be performed only when the computations of all the  $z_s^{k+1}$  are finished. Its computation is fairly intuitive, it averages the results of each branch over the control horizon (to lighten the notation we omitted the time bound  $t < L$  in (4d)). The ADMM dual variables are updated in (4e). Over the course of the optimization, these dual and consensus updates ensure that all  $z_s$  converge to a same consensus  $\tilde{z}$  for  $t < L$ .

In the literature, the ADMM algorithm usually refers to a decomposition into 2 subproblems that are solved and updated in a sequential fashion. We use here the ADMM variation called *consensus optimization* (see [17]) for a  $N$ -fold decomposition, and where the optimizations of the subproblems are parallelizable.

5) *Termination criterion*: The procedure can be stopped once the constraints of each subproblem are satisfied and once a consensus is reached for  $t < L$ . Formally it means that both the AuLa and ADMM primal and dual residuals are smaller than threshold values ( $\epsilon^{pri}, \epsilon^{dual}, \xi^{pri}, \xi^{dual} \in \mathbb{R}$ ).

$$\|g_s(z^{k+1})\| \leq \epsilon^{pri}, \|h_s(z^{k+1})\| \leq \epsilon^{pri}, \quad (5a)$$

$$\|z_s^{k+1} - z_s^k\| \leq \epsilon^{dual}, \quad (5b)$$

$$\|z_s^{k+1} - \tilde{z}_s^{k+1}\| \leq \xi^{pri}, \quad (5c)$$

$$\|\tilde{z}_s^{k+1} - \tilde{z}_s^k\| \leq \xi^{dual}. \quad (5d)$$

In the examples of the experimental section, the algorithm typically converges after 10 to 30 iterations.

## V. EXPERIMENTS

The solver is implemented in C++. The source code and a supplementary video are available for reference\*.

### A. Adaptive Cruise Control among pedestrians

We consider the problem briefly introduced in Figure 1. The car drives along a street in presence of pedestrians on the sides who may cross. The pedestrians' intentions are partially observed through a simulated perception module that outputs, for each pedestrian, the probability that it will cross in front of the car. Eventually pedestrians either cross the street, or walk on the walkway making their intention fully observable.

We optimize the longitudinal acceleration of the car. The car dynamic is modeled as a linear system (6) and we consider quadratic cost with linear constraints. The problem is reduced to an optimization in control space only (see section III-E). (2) takes the form of  $N$  loosely-coupled QPs.

Performance are evaluated on randomized scenes in simulation. The car dynamic is simulated in Gazebo. Planning occurs at 10Hz with a 5s horizon and 4 steps per second.

We compare the results obtained with the control-tree optimization vs. a baseline. The baseline is a classical MPC approach considering one single hypothesis. To make sure that no collision happens with pedestrians, the hypothesis used is the worst case: as long as the pedestrian's intention is uncertain, the car plans to stop in front of it.

1) *System dynamic*: Let  $x$  be the longitudinal coordinate of the ego-vehicle along the road. The system dynamic is described by the linear system:

$$\begin{pmatrix} x_{t+1} \\ v_{t+1} \end{pmatrix} = \begin{pmatrix} 1 & dt \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_t \\ v_t \end{pmatrix} + \begin{pmatrix} 0 \\ dt \end{pmatrix} (u_t), \quad (6)$$

where  $v_t$  and  $u_t$  are respectively the vehicle speed and the controlled acceleration.

2) *Partially observable discrete state*: With  $N_p$  pedestrians located at  $x_0 < \dots < x_{N_p-1}$  ahead of the vehicle, the discrete state can be described by an integer  $s \in [0..N_p]$  indicating the closest pedestrian who crosses.  $s = N_p$  is the case where no pedestrian crosses.

We note  $p_i$  the output of the simulated prediction module giving the crossing probability of the  $i^{th}$  pedestrian. The  $s^{th}$  pedestrian is the closest crossing if it crosses, and, the pedestrians before him don't cross, such that:

$$p(s) = p_s \prod_{i=0}^{s-1} 1 - p_i. \quad (7)$$

3) *Costs and constraints*: The control-tree is optimized w.r.t the following trajectory costs:

- **Speed**: The matrix  $Q = \begin{pmatrix} 0 & 0 \\ 0 & k_v^\dagger \end{pmatrix}$  penalizes the velocity difference between the vehicle speed and a given desired velocity.
- **Acceleration**: The square acceleration is penalized,  $R = (k_u^\dagger)$ .

In addition, the following constraints are applied:

- **Stop before the  $i^{th}$  pedestrian**: A state inequality constraint applies to stop and keep a safety distance to the pedestrian,  $x \leq x_i - d_{safety}^\dagger$ .
- **Control bounds**: Longitudinal acceleration is constrained to stay between bounds  $[-8.0, 2.0]m/s^2$ .

4) *Example of control-trees*: Figure 5 shows a control-tree obtained with a vehicle launched at 48 km/h (30mph) in the presence of 3 pedestrians. Each pedestrian has a probability 0.15 that it will cross. This implies a probability of 0.61 that the road is free. The control-tree doesn't brake too hard, but still guarantees that it is possible to come to a stop in the worst case (see red curve in Figure 5). On the other hand, the single hypothesis approach brakes much stronger.

5) *Influence of the belief state*: The Figure 6 focuses on the control horizon ( $t \leq 1.0s$ ) and shows the influence of the crossing probability. When the probability is low, the planned control is more optimistic, whereas when this probability increases, the control policy becomes more conservative and tends to the limit-case (single hypothesis).

6) *Evaluation on random scenarios*: The algorithm is tested under various combinations of pedestrian density and pedestrian behavior (average crossing probability). Each run is performed over 30 minutes of simulated driving. We report in table I on the average costs (as defined by the matrices  $Q$  and  $R$ ) within the control horizon. To give a sense of the conservativeness of the car, we indicate the average velocity.

In table I, we also compare the performance of two variations of the optimization versus the baseline: tree-2 and tree-5 have respectively 2 and 5 branches. With 80 pedestrians per km, up to 4 pedestrians can enter the planning horizon such that 5 branches are in principle, needed. Tree-2 is therefore an approximation that we use to evaluate the benefit of having larger trees versus the computation time.

\*<https://github.com/ControlTrees/icra2021>

<sup>†</sup>Table I is obtained with  $k_u = 5.0$ ,  $k_v = 1.0$ ,  $d_{safety} = 2.5$  m

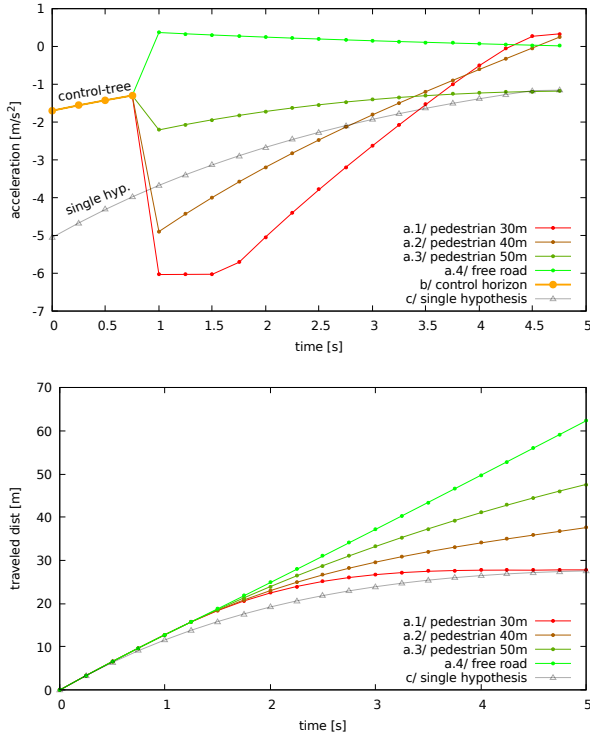


Fig. 5: Example of control-tree when braking: a.1, a.2, and a.3 anticipate that a pedestrian crosses. a.4 corresponds to the free road scenario. c. is obtained with single hypothesis MPC assuming that the closest pedestrian crosses.

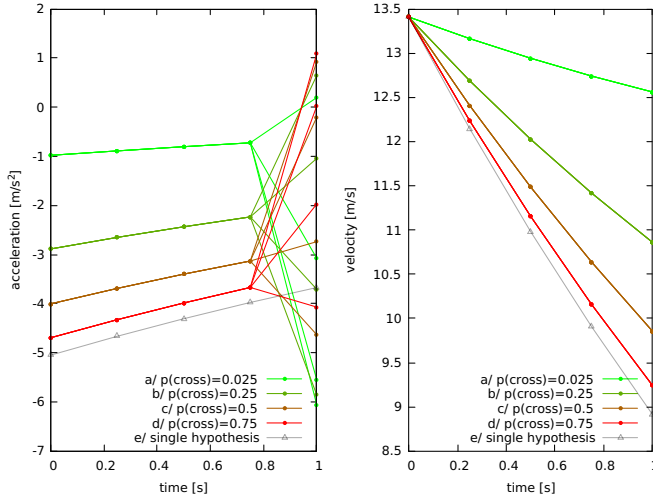


Fig. 6: Influence of the belief state on the braking within the control horizon: Low crossing probabilities (e.g. a, b) lead to a more optimistic control-tree.

7) *Results interpretation:* Control-trees always lead to better control costs than the single hypothesis MPC (up to twice lower costs in the case of 20 pedestrians per km, and 5% of crossing probability). In particular, the car drives less conservatively, maintains on average a higher velocity but still ensures that it is possible to come to a stop safely might a pedestrian cross. The benefit is more significant when the

	pedestr. per km	percentage of pedestr. crossing	avg. cost	avg. speed (m/s)	planning time (ms)
tree-2	20	5%	28.8	10.8	11.1
single hyp.	20	5%	60.3	8.15	4.92
tree-2	20	25%	69.3	7.68	17.8
single hyp.	20	25%	83.0	6.18	7.03
tree-5	80	1%	44.5	8.84	20.4
tree-2	80	1%	48.8	8.50	9.05
single hyp.	80	1%	86.2	5.45	3.05
tree-5	80	5%	63.9	7.31	20.1
tree-2	80	5%	68.3	6.96	9.61
single hyp.	80	5%	95.8	4.66	3.06
tree-5	80	25%	109.8	4.18	19.9
tree-2	80	25%	112.7	3.87	12.1
single hyp.	80	25%	119.5	3.24	5.52

TABLE I: Performance comparison: The benefit of control-trees is greater when the likelihood of the worst case is low.

crossing probability is low. This can be understood easily: by discarding the probabilities, the baseline is comparatively more conservative when the probability that the worst case happens is low.

Control-trees with 5 branches perform better than those with 2 branches only. This is consistent since the problem structure is more completely modeled with 5 branches. However, the benefit between tree-2 and tree-5 is rather small suggesting that, in this example, a simple tree is a good trade-off between performance benefit and computation time.

Control-trees take longer to optimize, but computations stay fast and compatible with a real-time application.

8) *Scalability:* We tested the optimization with up to 100 branches (see Figure 7). The *Distributed Augmented Lagrangian* algorithm is compared to an approach where the optimization is not decomposed but cast into one large QP, and optimized with an off-the-shelf solver (OSQP see [18]). Tests were conducted on a Intel® Core™ i5-8300H.




Fig. 7: Scalability: In contrast to an undecomposed optimization (red), the proposed algorithm (green) scales linearly.

For small number of branches, our solver is slower. This reflects the fact that it is not yet, as optimized as standard solvers. However, for higher number of branches, the algo-



rithm is drastically more efficient than the undecomposed optimization. In practice, it scales linearly with the number of hypotheses. Indeed, the number of subproblems increases with the number of hypotheses (since they are equal), but the size of each subproblem doesn't change.

### B. Slalom among uncertain obstacles

In this example, the ego-vehicle drives along a straight reference trajectory. Obstacles appear randomly. Obstacle detection is imperfect: there are *false-positives* i.e. obstacles are detected although they don't exist. This captures a common problem when working with sensors like radars that can suffer from a high rate of false detections. The simulated detection module also outputs an existence probability for each detected object. The closer the car gets to the obstacle, the more reliable are the observations. Once the distance to obstacle becomes lower than a threshold (randomized in the simulation), the object becomes fully observed: it gets a probability of 1.0 or disappears.

Optimization is performed in configuration space using the *K-Order-Motion-Optimization* formulation (KOMO, see [19][20]). The horizon is set to 5 s with 4 steps per second.

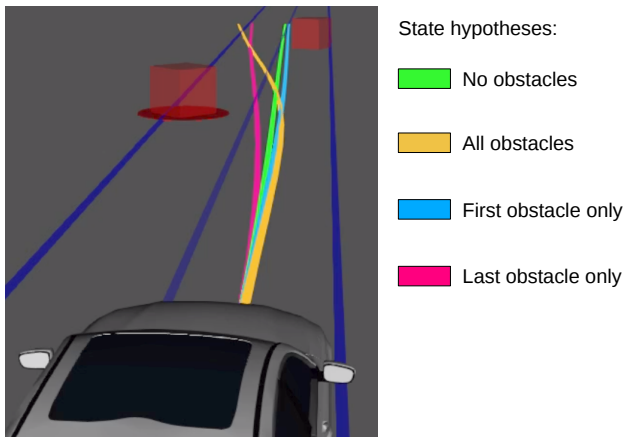


Fig. 8: Avoidance of 2 uncertain obstacles: Each trajectory corresponds to a different combination of object presences.

1) *State space and kinematic*: Optimization is performed in  $SE(2)$ . No slippage is assumed, this is enforced by the following non-holonomic constraint:

$$\dot{x}(t)\cos(\theta) - \dot{y}(t)\sin(\theta) = 0, \forall t, \quad (8a)$$

where  $(x, y, \theta)$  is the vehicle pose.

2) *Partially observable discrete state*: The discrete states represent the possible combinations of obstacle existences. With 2 uncertain obstacles in the planning horizon, there are 4 possible states (see Figure 8).

3) *Costs and constraints*: The following trajectory cost terms are minimized:

- **Acceleration**: The square accelerations  $\ddot{x}$  and  $\ddot{y}$  as well as the square angular acceleration  $\ddot{\theta}$ .
- **Distance to centerline**: The square distance to a reference line.
- **Speed**: The square difference to a desired velocity.

In addition, the following constraints apply:

- **Kinematic**: The non-holonomic constraint (see (8a)).
- **Collision Avoidance**: The distance to obstacle must stay greater than a safety distance.

Unlike the previous example, the non-holonomic and collision avoidance constraints make the problem *non-convex*.

Function gradients are computed analytically while we use the gauss-newton approximation of the 2<sup>nd</sup> order derivatives.

4) *Evaluation on random scenarios*: Table II gathers results obtained when simulating 30 minutes of driving, with on average, one potential obstacle every 17 meters resulting overall in 900 uncertain obstacles being encountered. As in the previous example, the baseline used here, is a single-hypothesis MPC which doesn't use the belief state information: all uncertain obstacles are avoided the same way, regardless of their existence probabilities.

	percentage of false positive	avg. cost	planning time (ms)
tree-4	90%	0.42	59.3
single hyp.	90%	0.91	27.7
tree-4	75%	0.97	74.2
single hyp.	75%	1.8	32.8

TABLE II: Performance comparison: Control-trees lead to lower costs.

Compared to the baseline, the control-tree consistently provides lower trajectory costs within the control horizon (approximately twice lower). In particular, the vehicle experiences less lateral acceleration while still avoiding successfully all real obstacles.

It is also interesting to compare those costs to the ideal fully observable case i.e. where the agent has perfect information and knows which obstacles are real, and which are false detections. This ideal case gives a lower bound of the trajectory costs which amount to 0.3 and 0.78 for the scenarios with 90% and 75% of false detections respectively. With this point of comparison, the benefits compared to the baseline is even more significant.

Computation time increases but remains largely compatible with the planning frequency of 10Hz.

## VI. CONCLUSION AND FUTURE WORK

We proposed a new MPC approach for problems where discrete and critical aspects of the world are partially observed. Optimization is performed using probabilistic information, which results in policies that outperform policies optimized under a single hypothesis MPC scheme.

An optimization algorithm, *Distributed Augmented Lagrangian* is proposed which leverages the decomposable structure of the optimization problem to improve scalability.

In future work, our ambition is to scale the approach to scenarios requiring an order of magnitude more states. As this may not be possible to do so in real time, we would like to explore ways to use this multi-hypotheses MPC framework to compute offline, expert demonstrations that can be used to learn control policies.

## REFERENCES

- [1] C. Phiquepal and M. Toussaint, "Combined Task and Motion Planning under Partial Observability: An Optimization-Based Approach," pp. 9000–9006, 05 2019.
- [2] A. Mesbah, "Stochastic Model Predictive Control: An Overview and Perspectives for Future Research," *IEEE Control Systems*, vol. 36, pp. 30–44, 12 2016.
- [3] T. A. N. Heirung, J. A. Paulson, J. O'Leary, and A. Mesbah, "Stochastic model predictive control - how does it work?," *Computers and Chemical Engineering*, vol. 114, pp. 158–170, 2017.
- [4] P. J. Campo and M. Morari, "Robust Model Predictive Control," *1987 American Control Conference*, pp. 1021–1026, 1987.
- [5] J. Lfberg, "Min-Max Approaches to Robust Model Predictive Control," *Linkping Studies in Science and Technology. Dissertations. No.*, 01 2003.
- [6] S. Lucia, T. Finkler, D. Basak, and S. Engell, "A new robust nmpc scheme and its application to a semi-batch reactor example\*," *IFAC Proceedings Volumes*, vol. 45, no. 15, pp. 69 – 74, 2012. 8th IFAC Symposium on Advanced Control of Chemical Processes.
- [7] S. Lucia, T. Finkler, and S. Engell, "Multi-stage nonlinear model predictive control applied to a semi-batch polymerization reactor under uncertainty," *Journal of Process Control*, vol. 23, no. 9, pp. 1306 – 1319, 2013.
- [8] S. Lucia, J. Andersson, H. Brandt, M. Diehl, and S. Engell, "Handling uncertainty in economic nonlinear model predictive control: A comparative case study," *Journal of Process Control*, vol. 24, 08 2014.
- [9] D. Kouzoupis, E. Klintberg, M. Diehl, and S. Gros, "A dual Newton strategy for scenario decomposition in robust multistage MPC," *International Journal of Robust and Nonlinear Control*, 12 2017.
- [10] C. Leidereiter, A. Potschka, and H. Bock, "Dual decomposition for QPs in scenario tree NMPC," pp. 1608–1613, 07 2015.
- [11] S. Subramanian, S. Nazari, M. A. Alvi, and S. Engell, "Robust nmpc schemes for the control of mobile robots in the presence of dynamic obstacles," in *2018 23rd International Conference on Methods Models in Automation Robotics (MMAR)*, pp. 768–773, 2018.
- [12] E. Klintberg and S. Gros, "A parallelizable interior point method for two-stage robust mpc," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 6, pp. 2087–2097, 2017.
- [13] M. L. Littman, A. R. Cassandra, and L. P. Kaelbling, "Learning policies for partially observable environments: Scaling up," 1995.
- [14] M. Diehl, *Optimization Algorithms for Model Predictive Control*, pp. 1–11. London: Springer London, 2013.
- [15] M. Toussaint, "A novel augmented lagrangian approach for inequalities and convergent any-time non-central updates," 2014.
- [16] A. Conn, N. 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, 04 1991.
- [17] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, p. 1122, Jan. 2011.
- [18] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: An Operator Splitting Solver for Quadratic Programs," in *2018 UKACC 12th International Conference on Control (CONTROL)*, pp. 339–339, Sep. 2018.
- [19] M. Toussaint, "Newton methods for k-order Markov Constrained Motion Problems," 2014.
- [20] M. Toussaint, *A Tutorial on Newton Methods for Constrained Trajectory Optimization and Relations to SLAM, Gaussian Process Smoothing, Optimal Control, and Probabilistic Inference*, pp. 361–392. Cham: Springer International Publishing, 2017.