# BACHELOR'S DEGREE THESIS

Degree in Mathematics

Title: A Learning Approach To The FOM Problem

Author: Eudald Romo Grau

Advisor: Maria Alberich

Department: Matemàtiques

Academic year: 2016-2017

UNIVERSITAT POLITÈCNICA DE CATALUNYA
BARCELONATECH
UPC
Facultat de Matemàtiques i Estadística

Universitat Politècnica de Catalunya

Facultat de Matemàtiques i Estadística

Degree in Mathematics
Bachelor's Degree Thesis

A Learning Approach
To The FOM Problem

Eudald Romo Grau

Supervised by Alberto Rodriguez Garcia and Maria Alberich Carramiñana

April, 2017

Abstract

The family of modes (FOM) approach to solving model predictive control (MPC) problems is a novel heuristic technique developed at MIT to reduce the exponential time complexity of traditional MPC solving techniques. This study strives to bring this method closer to the manipulation research community by addressing some of the weaknesses of FOM, by comparing its results to traditional MPC solving techniques and by providing methodologies to choose the parameters required for the heuristic, that will ultimately lead to a FOM-based algorithm capable of tracking complex trajectories.

Keywords

1. Introduction

Hogan recently provided an heuristic technique called family of modes (FOM) [16] to respond to the current need for the robotics manipulation community to solve model predictive control (MPC) problems under hybrid constraints and underactuation fast enough to be introduced into control loops. The goal of our study is to bring the FOM technique closer to the manipulation research community.

In order to do that, we reinforce the FOM method by fixing some of its weaknesses -as its sequential properties-, we provide comparison tools to test the method against traditional hybrid MPC solving techniques -as mixed integer quadratic programming (MIQP)-, and we provide simple and systematic techniques to set-up the method's parameters that will ultimately leat to a new variation of the FOM approach that can properly track complex trajectories.

We also discuss the new lines of research we plan on exploring before presenting this study to 2017's International Symposium on Robotics Research (ISRR). They can be summarized as:

- To proof the scalability of this method by increasing the complexity of the considered systems beyond the one used in Hogan's work, that tracked straight lines with a single point pusher and a square slider.

- To check the robustness of the FOM method by using it in models different from those it was designed for.

We hope that the results from our studies will increase the performance and simplify the implementation of FOM and will bring it closer to the robotics manipulation research community by both providing systematic set-up tools and showing its robustness and scalability.

To present our studies and results, we start by providing in chapter 2 all the required theoretical background in planar pushing manipulation, that will allow ultimately lead to the concept of hybridness, a concept fundamental to understand the behaviour, the difficulty and the motivation of the FOM controller. Chapter 3 is completely dedicated to familiarize the reader with hybrid control.

Once the theoretical background is presented, the goal of this study is revisited in chapter 4 emphasizing the implications on the manipulation research community of being able to solve trajectory-control MPC problems fast.

The necessary formulation is derived afterwards in chapter 5 and any theoretic concept that still needs further clarification will be re-discussed here.

Having presented the background, objectives and formulation of our study, we present the core of our research. We separate it into two main groups. In chapter 6 we study the properties of the FOM method and we test them in a simulation interface we developed. In chapter 7 we provide a procedure to select the modes of the FOM and we test it in the simulation interface of the previous chapter. The mode selection procedure will ultimately lead to the concept of Dynamic Family of Modes (DFOM), a trajectory optimization method based on FOM that is capable to track complex trajectories.

To conclude the study, all the results we obtained are discussed in chapter 8 together with their implications. This discussion ultimately leads to chapter 9, that presents the future lines of research we already mentioned at the beginning of this section.

2. Planar Pushing

As Mason explained in his talk during his CITRIS People and Robots Seminar Series talk [20], robotics manipulation research has been traditionally based on a pick and place methodology that consisted in solving two fundamental problems:

1. **Determine a stable grasping configuration:** A stable grasping configuration was generally obtained by finding a set of points that completely constrained the picked object position and orientation with respect to the gripper.

2. **Solving a collision-free trajectory planning problem:** Once the object was grasped it was rigidly attached to the robot gripper and the trajectory planning problem was completely determined at every point.

This methodology was constantly improved as the manipulation community considered broader and more complex problems, as Rodriguez' and Mason's study [26], that generalized the concept of stable grasp to caging -a partial immobilization of the grasped object that allows it to move in a bounded space. Other modifications added further constraints to the problem, as the concept of compliant grasping manipulation [21], where the object motion of the object is constrained by the task and the constraint is enforced without exerting heavy force loads on the object. An example of compliant manipulation the process of activating a lever. The lever only has one degree of freedom and can easily be broken by small planning uncertainties if a powerful robotic arm is used and the trajectory is controlled uniquely with the robot position without considering the applied forces. It should be noted that all these variations were applied to the same pick and place grasping problem. Some of the main motivations that made stable grasping so popular are:

- **Easy to track:** The obtained solutions are constrained by the contact points and the gripper pose so, even if the object could move inside the gripper (as in cagings) the object position could be deterministically bounded during all the trajectory.

- **Robustness:** Most grasping point predictions are robust to small sensor noise and applying enough grasping force allows ignoring uncertainties in the grasped object friction coefficients.

- **Simplicity:** As explained, the manipulation problem could be formulated simply by a grasping point identification and a path planning problem.

At least as far as on 1980's, some drift towards the study of the mechanics of planar pushing started to appear with studies as Mason's [22]. The main motivations for this drift were to handle problems that were not tractable with grasping or to better understand the grasping mechanics. A non-exhaustive list of the motivations of planar pushing could be:

- **You don't need to lift:** Lifting heavy objects using a traditional robotic arm can put heavy torque strain on the arm in certain arm configurations and the strain can break or topple the arm. If most or all of the planned trajectory can be done through pushing, the required force is usually reduced and when the maximum force is achieved safety mechanisms can be implemented so that the arm doesn't break.

- **Sometimes you need to push to grasp:** From the observation of human object handling, it can be seen that sometimes we prefer or need a previous pushing manipulation before grasping an object. As can be observed in Figure 4, constrained spaces like a shelf may force us to previously push a book against one of the self walls to pull it outside so that we can grasp it.

- **It Can Represent In-Hand Re-grasping:** Humans constantly show the ability and need to modify a grasping configuration without letting the object out of the grasp. When we write with a pen or drink from a bottle, we are constantly adapting our initial grasp without letting the object go to achieve more comfortable configurations. Some studies as Chavan's [6] [**?**] have formulated this in-hand re-grasping as a planar pushing problem by considering either one moving finger or an external surface as the pusher and the contact between the fixed fingers and the object is described as planar sliding.

- **Weaker Quasistatic Assumptions:** As will be further explained in section 5.2, in order to use quasistatic grasping analysis (commonly done in the manipulation community) a stronger set of assumptions has to be taken than to use quasistatic pushing.
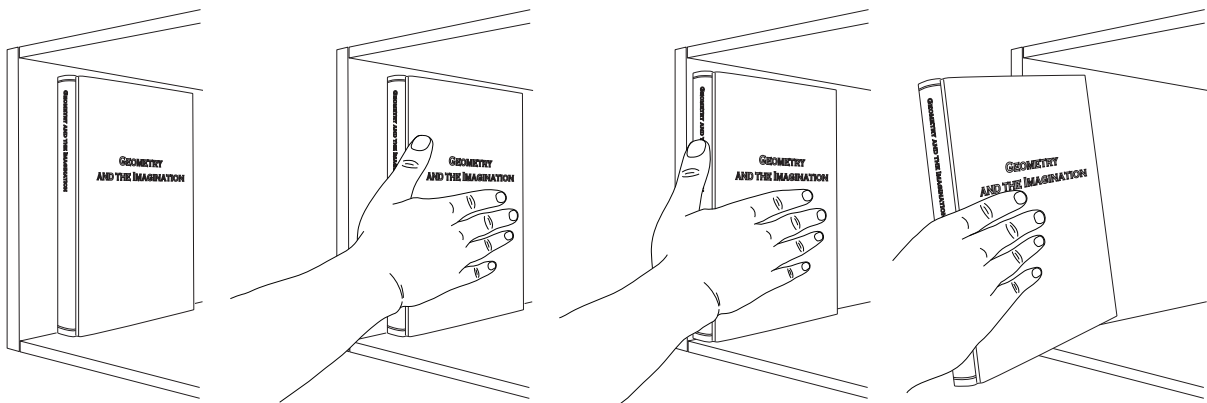

Figure 1: Figure from [16] that shows a human-like book grasping procedure.

The mechanics of planar pushing are complex and contain some inherent qualities as underactuation and hybridness that need to be addressed in order to solve planning, simulating or control problems. In the following sections we are going to introduce the basic assumptions taken and models used in planar pushing, together with the difficulties that arise from them.

2.1 Coulomb Friction

One of the usual assumptions in the robotics community is considering dry Coulomb contact forces. This model separates friction into two possible states (static and dynamic or kinetic friction) and the set of rules which describe it are a combination the work of Amontons [3] and Coulomb [8]:

- Amontons's first law: The force of friction is directly proportional to the applied load.

- Amontons's second law: The force of friction is independent of the apparent area of contact.

- Coulomb's law: Kinetic friction is independent of the sliding velocity.

This laws are usually condensed by defining the static (μ_s) and kinetic (μ_k) coefficients and considering the normal (f_n) and tangential (f_t) components of the contact force:

Theorem 2.1. Coulomb's Law:

$$f_t \begin{cases} < \mu_s f_n & \text{(for static contact)} \\ = \mu_k f_n & \text{(for kinetic contact)} \end{cases}$$

Where f_t is determined in the first situation by force balancing.

Remark 2.2. Through the rest of this thesis we use the common assumption taken in robotics manipulation of assuming a single coefficient of friction. Thus the previous theorem is simplified to $f_t \leq \mu f_n$

Mason cites prior statements by Da Vinci (from Truesdell's notes [29]) and experimental verification by Truesdell and Guillmor [11] and I would like to add the work of Euler [23], who first distinguished between static and kinetic friction.

2.2 Contact Slider-Support Surface

In planar pushing, there are two different contacts to be considered: the contact between the object (slider) and the support surface and the contact between pusher and slider. In this section we are going to introduce the main theoretical results about the first one. We will start with Mason's study of the relation between friction force and both motion and instantaneous center of rotation (which will be used again in section 5.2) and we will end with Goyal's Limit Surface, that will provide a geometric interpretation to this relation.

The contact between the slider and the support surface is inherently complex. If the slider is supported by more than three contact points the distribution of support normal forces is indeterminate and, if Coulomb contact is assumed, this indetermination is propagated to the friction forces. The results in this section are based on the previous assumption that the pressure distribution is known (for our study we will consider a homogeneous distribution).

Assuming a known pressure distribution and considering **r** the position vector of a point of the slider, **v(r)** its velocity vector, and p(**r**) the pressure at that point, then the total force and moment due to friction can be expressed as:

Definition 2.3. Total force due to friction:

$$\mathbf{f}_f = -\mu \int_A \frac{\mathbf{v(r)}}{|\mathbf{v(r)}|} p(\mathbf{r}) dA \tag{1}$$

Definition 2.4. Total moment due to friction:

$$\mathbf{n}_f = -\mu \int_A \mathbf{r} \times \frac{\mathbf{v(r)}}{|\mathbf{v(r)}|} p(\mathbf{r}) dA \tag{2}$$

If the slider trajectory is assumed to be a rotation over a instantaneous center $\mathbf{r}_{IC}$, the velocity can be expressed as $\mathbf{v(r)} = \omega \times (\mathbf{r} - \mathbf{r}_{IC}) =$ and the previous results can be rewritten as:

Definition 2.5. Total force due to friction:

$$\mathbf{f}_f = -\mu \, sgn(\dot{\theta}) \, \hat{\mathbf{k}} \times \int_A \frac{\mathbf{r} - \mathbf{r}_{IC}}{|\mathbf{r} - \mathbf{r}_{IC}|} p(\mathbf{r}) dA \tag{3}$$

Definition 2.6. Total moment due to friction:

$$n_{fz} = -\mu \, sign(\dot{\theta}) \int_A \mathbf{r} \frac{\mathbf{r} - \mathbf{r}_{IC}}{|\mathbf{r} - \mathbf{r}_{IC}|} p(\mathbf{r}) dA \tag{4}$$

Remark 2.7. These equations are well defined when the magnitude of $\mathbf{r}_{IC}$ tends to infinity and they correspond to pure translations perpendicular to it.

These equations provide a relation force-motion for known continuous pressure distributions but if non-zero support force is allowed at discrete points, they become undefined for rotations about each of them. In fact, the relation between motion and frictional force cannot generally be described as an analytical function. Goyal et al. [12] [13] applied concepts of classical plastic theory to describe planar sliding, deriving the concepts of limit surfaces for dry contact manipulation, that describe a mapping between motion and force.

Goyal's work can be generalised to friction models which are more general but, for clarity, the main concepts will be introduced with the special case of isotropic Coulomb friction (as was presented in their original paper).

Consider a rigid body sliding over a surface with a known or pressure distribution p($\mathbf{r}$) at every contact point and assume that each point of contact follows isotropic Coulomb friction law. Let **f(r)** be the differential frictional force applied by each slider particle to the surface. Coulomb law allows us to obtain two important properties:

1. **Limit Curve:** The set of possible point differential friction forces is contained inside a circle of radius $\mu p(\mathbf{r})$ centred at the origin of the force space. We will call this circle the limit curve (LC) of each contact point.

2. **Maximum Power Inequality.** The following relation is held at every contact point:

$$\forall \mathbf{f(r)}^* \in LC : \; (\mathbf{f(r)} - \mathbf{f(r)}^*) \, \mathbf{v(r)} \geq 0$$

   Where **v(r)** is the point velocity at each contact point and $\mathbf{f(r)}^*$ is an arbitrary force that can take any possible value inside the LC.

The maximum-power inequality and limit curve are concepts introduced from classical plastic theory that can be interpreted as follows: when slip occurs **v(r)** is non-zero and, for the case of a circular LC, the maximum power inequality restricts **f(r)** to have the same direction as **v(r)** and lie on the LC boundary. That is, the frictional load is on the limit curve and the motion vector is normal to the limit curve at **f(r)**. When the point velocity is zero, the frictional load is undetermined.

Goyal's study further uses classical results of plastic theory to generalize this concept to general closed LC that are convex (to allow the maximum power inequality to be held at every contact point) and that enclose the origin of the force space (to allow positive energy dissipation). The motion-force relationship properties obtained in this general LC study are summarized in Figure 3 and explained here:

1. If **f(r)** is in the interior the LC, then **v(r)** is zero.

2. If **f(r)** is on a LC boundary point with a well-defined normal, then **v(r)** is perpendicular to the LC.

3. If **f(r)** is on a LC vertex, the direction of **v(r)** is not determined.
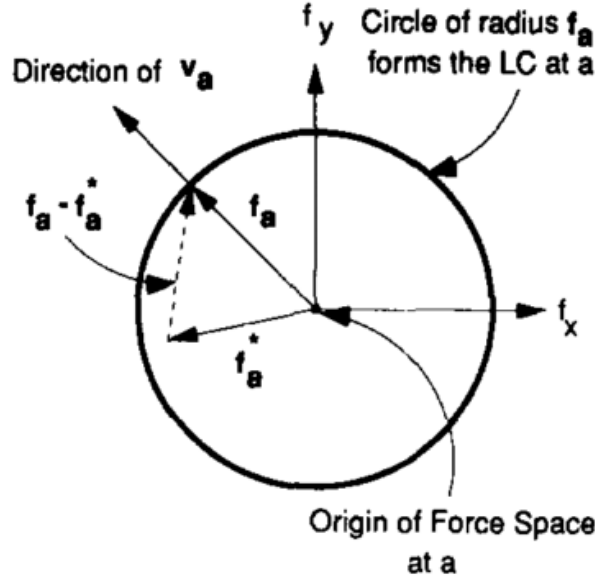
Figure 2: Figure from [12] that shows a limit curve for isotropic Coulomb contact at a certain position r_a. A unique set of v_a and f_a is determined when the point is slipping and it corresponds to the points in the circumference. When the force lies inside the circle, the point does not move and the frictional force is not determined.

4. For all $\mathbf{f(r)}$ in a same flat region of the LC, a common $\mathbf{v(r)}$ value is associated, perpendicular to the LC, so the direction of $\mathbf{f(r)}$ is not determined given $\mathbf{v(r)}$

If we consider $\mathbf{p}$ the total frictional load wrench

$$\mathbf{p} = \begin{pmatrix} f_x \\ f_y \\ n_{0z} \end{pmatrix} = \sum_r \begin{pmatrix} f_x(\mathbf{r}) \\ f_y(\mathbf{r}) \\ \mathbf{r} \times \mathbf{f(r)} \end{pmatrix}$$

and $\mathbf{q}$ the velocity twist

$$\mathbf{q} = \begin{pmatrix} v_{0x} \\ v_{0y} \\ \omega_z \end{pmatrix}$$

then Goyal's study generalizes the concept of single point Limit Curves for force-motion relationship to Limit Surfaces(LS) that relate total wrench and twist for the slider:

Theorem 2.8. Limit Surface:

1. *The set of possible load wrenches $\mathbf{p}$ is contained inside a LS.*

2. *The following relation is held at every contact point:*

$$\forall \mathbf{p}^* \in LS : \ (\mathbf{p} - \mathbf{p}^*)\mathbf{q} \geq 0 \tag{5}$$
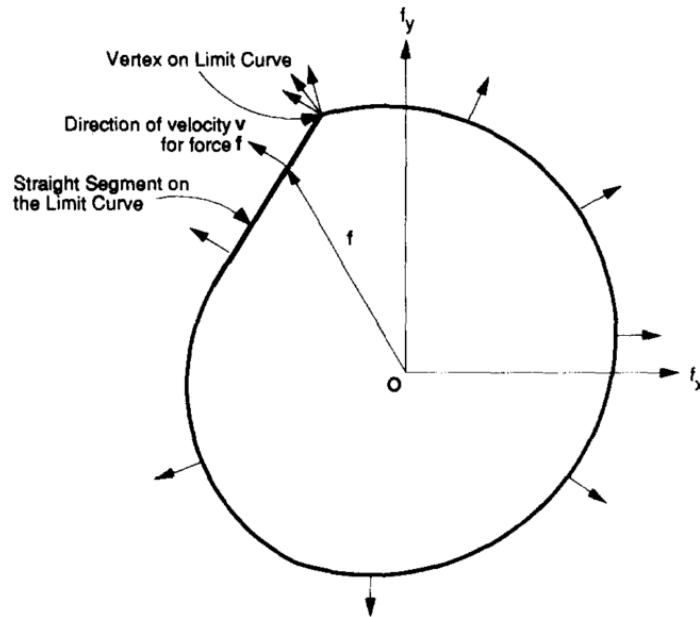
Figure 3: Figure from [12] that shows a general limit curve at a certain contact point. The four possible motion-force relations are represented on it.

Certain properties are derived for the LS and we are going to list here without proof the ones that apply to our study:

- With isotropic friction, flat regions appear on the limit surface if and only if it has points of finite support (infinite stress).

- With isotropic friction, vertices appear on the limit surface if and only if all support points lie on one straight line segment.

When smooth enough, the limit surface allows to compute the frictional wrench from a non-zero motion twist, but it should be noted that the reciprocal relation can never be obtained (except for the trivial case where the wrench is inside the surface and the twist is zero), but the twist direction can be obtained successfully. As will be discussed in chapter 5, combining the constraints obtained from this relation with other constraints applied on the Pusher-Slider contact, the motion equations for the pusher slider can be derived.

2.3 Contact Pusher-Slider. Friction Cone And Hybridness

The contact between the pusher fingers and the slider sides could be described as in the previous section (by considering the finger a slider and the slider side a support surface) but the nature of the problem allows some simplifications. The fingers we consider never rotate with respect to the slider side, so all the motion is pure translation. Furthermore, we only consider planar pushing (we do not consider operations as tilting the slider), so we are only interested in moving the pusher along one dimension of the slider face (for most purposes we consider the slider to be flat).

A common and useful geometric interpretation of Coulomb's law called the friction cone (according to Mason first constructed by Moseley [24]) is widely used in the robotics community for one dimensional single point contacts. Given a point on a surface interacting with it with total contact force **f**, we decompose it into its normal f_n and tangential f_t components with respect to the contact surface. Coulomb friction compact formulation states that $f_t \leq \mu f_n$, where μ is the proportionality factor usually called friction coefficient, where the equality holds on dynamic friction. Letting α be the angle between **f** and f_n (or, alternatively, the normal of the surface on the point of contact), Coulomb's law is equivalent to:

Theorem 2.9. The set of feasible forces in dry Coulomb contact must satisfy:

$$\alpha := tan^{-1}\frac{f_t}{f_n} \leq \mu$$

So, the set of feasible contact forces lies inside a closed cone with aperture 2μ.

Figure 4:   Figure from [22] that shows the friction cone, the geometric interpretation of Coulomb's Law. The set of all vectors following.

Consider a single point finger pushing a planar slider. There are three possible contact states: The finger can stick to the slider (Sticking State), the slider can slide with a slip to the left relative to the finger (Sliding Left State) or it can slide with slip to the right (Sliding Right State). This section shows that in each state the pusher can apply a different range of forces to the slider. Each of this different range of forces will translate into different ranges of friction reactive wrench, that will ultimately lead to different slider twist ranges.

This multiple system state phenomenon (inherent in Coulomb Contact) is commonly referred as hybridness. Hybridness appears in many manipulation situations different than sticking/sliding ones. For example a gripper can throw the object in the air to re-grasp it, change the number of fingers contacting the object, change the pusher-slider contact side while pushing a slider of irregular shape, etc. It also appears in other robotic fields as locomotion, where there are always at least two differentiated states (aerial state and ground state) with clearly different dynamics.

In the following chapter, we are going to discuss the hybridness phenomenon in more detail by covering why is it important to take into account, the problems that arise when considering hybridness in planning and control problems and different methods to solve them in each field, that will ultimately lead to the introduction to the FOM method.

3. Hybridness in Manipulation Problems

Through all this section, we will compare the role of hybridness in different aspects of manipulation problems. The principal fields we will distinguish are simulation, trajectory planning and trajectory control. Each of these fields take a different role when studying a general manipulation problem, that we can generally defined as follows:

Definition 3.1. Manipulation Problem:

Given a system conformed by a set of points in a metric space M and a parametrization $c(\mathbf{x})$ (for $\mathbf{x}$ in a certain configuration space C) that uniquely describes the position of all the points of interest in the metric space, a manipulation problem consists in changing a certain initial configuration $\mathbf{x}_0$ into a final configuration $\mathbf{x}_f$ following appropriate constraints.

Simulation problems consist on accurately reproducing the real physics of a manipulation problem to safely and quickly study its properties. A key property of most simulation problems is that, at every instant of time, we know the configuration of the system until the current time step and the only variables we need to compute are the ones that define the configuration in the next step.

Planning problems consist on finding a trajectory in the configuration state to go from state $\mathbf{x}_0$ to $\mathbf{x}_f$. Most simple planning problems have a state variable that defines the configuration of the object we need to manipulate and a actuator state variable that describes the configuration of the gripper or other tool used to move the object. As opposed to simulation problems, planning problems don't have full information about the state of the object and actuator state variables for each step of the trajectory, and they are usually variables of the problem. Because of this difference, the equations that describe the evolution of the same actuator-object system can have different properties in planning or simulation problems. As an example, the motion equations we derive in section 5.5 are generally non-linear (they are linear with respect to the object state variables but not with respect to the actuator state variables), but they can be expressed as linear equations on the actuator state variables for the first step of the problem, where we know the initial state of the object.

Traditional control designs as PID controllers can successfully control a system by solving only local manipulation problem, in a similar way as simulation. PID only use the notion of the current and past differences between the configuration of the object $\mathbf{x}$ and its objective value $\mathbf{x}^*$ to devise a control response to stabilize the realized trajectory around the objective one. Such control cannot be used for pushing because of a phenomenon called underactuation.

3.1 Underactuation. Trajectory Optimization

A control problem with first order motion constraints $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$ for some control input $\mathbf{u}$ is underactuated when f is non-exhaustive in the configuration velocities space. In most problems, $\dot{\mathbf{x}}$ is the result of converting a second order Newtonian equation $\ddot{\mathbf{x}} = m \sum \mathbf{F}$ to a system of first order ordinary differential equations (ODEs) and underactuation is usually caused by the controller being able to transmit only a limited set of forces and moments to the object.

In stable grasping, the system is fully activated (see remark 3.2) at all times. On the other hand, for a single simple point finger pushing the system is highly underactuated, as the normal force the pusher applies to the slider is non-negative (there cannot be no pulling forces) and the overall contact force lies in the friction cone. As such, the controller must reason only among the forces that can physically be realized at every instant.

In a certain way, it can be considered that underactuation distorts the usual metric of the system. The example in Figure 5 shows a tracking problem where a single point pusher needs to move a square slider to a configuration situated just behind the current pusher location. Any local pushing control action will bring the system away from the objective configuration if we use a tradition 2D Euclidean metric. If the controller takes into account a certain notion of path and future and tries to minimize the overall distance through the considered control trajectory, it can obtain a solution like the one we show.

Remark 3.2. Stable grasping has some some sources of underactuation, but they are linked with the situations were the stable grasp property cannot be held. For example, if an object is too heavy to be lifted or to be kept in a grasp without slipping. These are examples of underactuation, as the gripper cannot transmit enough force to the object, but they are also examples where stable grasping is not possible.

A common method to consider future paths in control problems is to solve simplified planning problems at high frequency. The planning problems can usually be expressed as trajectory optimization problems, where the values of the state and control variables are chosen to minimize a cost function along the trajectory, while fulfilling the required motion constraints that represent the kinematics of the real world problem. These optimization problems are usually linearized around the nominal trajectory computed by a planning problem beforehand. This linearization reduces the computational cost of the control problem while ensuring that, in a certain region of attraction, the controlled trajectory converges to the nominal one. It is also common to only consider the control path until a certain finite time horizon (as opposed to planning problems, that usually consider the trajectory until reaching a certain goal).

Once stated the need to consider trajectories in pushing control problems we will now justify the need to consider hybridness in the control problems. This justification is threefold:

- **Optimality:** If we consider all the possible modes in a trajectory optimization problem, there's a wider set of possible feasible states and control laws, so we may be able to obtain a better solution.

- **Robustness:** We need to consider the possible hybrid states and their dynamics so that we can react to execution errors or model uncertainties. We need to understand the mechanics of slipping to be able to return to our optimal pushing point after an accidental slip.

- **Stability:** In section 6.2 we will show that when controlling a hybrid system through its perturbation motion equations, if the nominal trajectory is not feasible in our formulation we will not be able to stabilize our control trajectory around it, because the nominal trajectory will not even be an equilibrium point of our motion equations.

Figure 5: Control trajectory obtained when trying to bring the slider to a configuration behind the current position of the pusher. The red slider describes the objective configuration, the blue line describes the trajectory of the slider center of mass during the control trajectory and the blue sliders represent some of the configurations of the system during the tracking trajectory. The position of the slider when we stopped the simulation is highlighted with thicker and more opaque lines.

In the next sections we will introduce formulation techniques to add the hybridness concept into manipulation problems. In particular, we will discuss the limitations of some of the formulations to solve control problems at enough frequency. This will ultimately lead to Hogan's FOM formulation as an approximate method to solve the finite horizon trajectory optimization problem.

3.2 Complementarity Problems

Given a variable $\mathbf{z}$ and a function on it $\mathbf{g(z)}$, a general complementarity problem consists on finding a value for $\mathbf{z}$ that satisfies:

$$\mathbf{z} \geq \mathbf{0} \qquad \mathbf{g(z)} \geq \mathbf{0} \qquad \mathbf{g(z)}\,\mathbf{z} = 0$$

This formulation derives from traditional linear programming techniques to encode two mutually exclusive constraints on $\mathbf{z}$ ($\mathbf{z} \geq \mathbf{0}$ and $\mathbf{g(z)} \geq \mathbf{0}$). The linear complementarity problem (LCP) [7] assumes a linear g function and takes the following form:

$$\boldsymbol{\omega} = \mathbf{Rz} + \mathbf{u}$$

$$\boldsymbol{\omega} \geq \mathbf{0} \qquad \mathbf{z} \geq \mathbf{0} \qquad \boldsymbol{\omega}^T \mathbf{z} = 0$$

It was first introduced to predict instantaneous acceleration of rigid bodies in 1996 by Pang and Trinkle [25] and Anitescu and Potra [4], that guaranteed a solution for this formulation on multi-rigid-body Coulomb friction contact and impact problems. This provided a systematic way to formulate hybrid constraints for simulation purposes. Soon after the instantaneous formulation introduction, Steward and Trinkle [27] proposed a time-stepping formulation in the space of impulses and velocities which increased the stability performance of LCP based forward simulation techniques.

Remark 3.3. It should be noted that the LCP formulation of the multi-rigid-body Coulomb friction problems used a polyhedral approximation to the friction cone.

There are fast algorithms to solve LCP problems that allow the manipulation community to use them in simulation, as in Steward and Trinkle's study [27], or other instantaneous-nature problems as the studies of Fazeli et al. to identify friction parameters from observed object motions [10] [9].


Figure 6: Figure from [27] with the simulation results obtained when emulating a falling and spinning rod. The simulation step was h = 0.0025 and the simulation time at each time instant is shown in seconds on one of the bar ends.

In instantaneous Coulomb contact problems the contact dynamics can be expressed in linear discrete

form through Steward and Trinkle's formulation, but many Newtonian motion equations take a non-linear form in trajectory optimization problems. Then problem becomes non-linear (NCP) and Lemke's algorithm, the key idea used to provide the LCP solution, cannot be used. The complementarity constraints are usually non-convex and can only be solved in a fast way by using the previously mentioned algorithm. In the case of NCP, it has to be solved with a commercial optimizer and the problem usually cannot be solved fast enough for neither planning nor control.

Certain trajectory optimization techniques, as Rapidly expanding Random Trees (RRT) can function solely considering instantaneous dynamics. In these cases LCP can not be generally used as is to solve the trajectory optimization problems, because they also contain non-complementarity subject variables that must hold linear constraints and LCP does not suppor them. Luckily, these problems, called mixed LCP (MLCP) can also be solved with a generalization of Lemke's algorithm. Nikhil's work on extrinsic dexterity [6] is an example of MLCP usage in planning. Unfortunately these trajectory optimization techniques used in combination with MLCP are too slow to use in most control problems.

3.3 MPC and MIQP

MLCP is usually too slow to be used in trajectory optimization based control problems. Complementarity constraints are generally non-convex, which makes it difficult to use them in other algorithms than LCP or MLCP problems. To tackle the computing cost of complementarity constraints, model predictive control(MPC) can be used. Given an initial state $\bar{\mathbf{x}}_0$, an MPC problem consists in defining a discrete set of N positions $(\bar{\mathbf{x}}_1, ..., \bar{\mathbf{x}}_N)$ and control inputs $(\bar{\mathbf{u}}_0, ..., \bar{\mathbf{u}}_{N-1})$ to define a trajectory a trajectory. $\bar{\mathbf{x}}_i = \bar{\mathbf{x}}(t_j)$, where $t_j = t_0 + hj$ for a certain time step h. Similarly, $\bar{\mathbf{u}}_i$ stands for the control input necessary to transition from $\bar{\mathbf{x}}_i$ to $\bar{\mathbf{x}}_{i+1}$. Given this notation, MPC optimizes the values of all $\bar{\mathbf{x}}_i$ and $\bar{\mathbf{u}}_i$ with respect to a certain cost function:

$$J(\bar{\mathbf{x}}_0, ..., \bar{\mathbf{x}}_N, \bar{\mathbf{u}}_0, ..., \bar{\mathbf{u}}_{N-1}) = \bar{\mathbf{x}}_N^T \mathbf{Q}_N \bar{\mathbf{x}}_N + \sum_{i=0}^{N-1} \left(\bar{\mathbf{x}}_i^T \mathbf{Q} \bar{\mathbf{x}}_i + \bar{\mathbf{u}}_i^T \mathbf{R} \bar{\mathbf{u}}_i \right) \tag{6}$$

Where terms $\mathbf{Q}$, $\mathbf{Q}_N$, and $\mathbf{R}$ are weights matrices that penalize the error state, final error state, and control input, respectively. MPC problems consider the hybrid state to be constant during state state transitions. Under this assumption, a hybrid mode can be defined:

Definition 3.4. A **hybrid mode** is a ordered set of hybrid states.

In MPC, an N-length mode completely fixes the hybrid states the system is in along the finite horizon considered trajectory.

Remark 3.5. We use the $\bar{(.)}$ notation because, as will be explained in sections 5.5 and 5.6, we will use the perturbation of the state and control variables with respect to a nominal trajectory and control inputs as the MPC variables.

Remark 3.6. MPC problems are based in the linear quadratic regulators (LQR) controllers. These controllers are used for unconstrained non-hybrid problems described by the motion equation $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$ and use the same cost functions as MPC. LQR are able to give an analytical control law to minimize the cost function provided for infinite horizon problems. Their term $\mathbf{Q}_f$ can be set-up to account for the integral of the cost function from the last point considered until reaching the destination (if the trajectory is finite) or until infinite. If the $\mathbf{Q}_f$ term is chosen to follow the Riccati equation:

$$\mathbf{Q}_f = \mathbf{A}^T \mathbf{Q}_f \mathbf{A} - \mathbf{A}^T \mathbf{Q}_f \mathbf{B} \left(\mathbf{B}^T \mathbf{Q}_f \mathbf{B} + \mathbf{R} \right)^{-1} \mathbf{B}^T \mathbf{Q}_f \mathbf{A} + \mathbf{Q}$$

then the finite horizon LQR controllers will behave as the infinite horizon ones. In MPC problems there's no closed expression for this term but there are certain techniques to choose a good value to maximize the final region of attraction as the one exposed in Ya-Feng's et al. study [30], but it is generally based on the LQR technique of solving the Riccati equation . In our model we use a simplification of Ya-Feng's approach which simply computes the solution of the Riccati equation presented above using the motion equation of the last MPC step, that can be expressed as $\dot{\mathbf{x}} = \mathbf{A}_N \mathbf{x} + \mathbf{B}_N \mathbf{u}$.

If we consider a manipulation problem with convex dynamics, constraints, and cost function, the overall optimization problem is convex if a hybrid mode that describes the finite horizon is provided. Thus, the global problem consists in selecting the optimal solution from the finite set of optimal solutions for each of the fixed-mode sub-problems.

The number of possible modes of length N is h^N, where h is the number of possible hybrid states. This exponential growth generally makes a naive exploration of all the possible sub-problems too computationally expensive. In order to reduce the algorithmic strain, a mixed integer programming (MIP) strategy is usually adopted. This will be further explained in chapter 5, but the main idea is that the state space is expanded by adding h binary variables per MPC step that define whether the step i of the MPC is in a hybrid state s. The commercial optimizers provide tools to express state the state related constraints with respect to each of this binary variables so that they only apply when their correspondent variable is activated. Then, if possible, the optimizer will use internal MIP techniques as branch and bound and cutting planes [1] to avoid exploring all the possible modes. In our particular case, we will be able to express all the motion and friction constraints as linear inequalities and the objective cost is a quadratic function on the variables. This kind of mixed integer problems are called mixed integer quadratic problems (MIQP) and they have very fast commercial optimizers (as the one we will use, called Gurobi).

The drawback of MIQP is that it doesn't escalate well with the number of possible hybrid states and, even with a small number of states, the optimizers aren't generally fast enough to solve the problem online, as required in control problems. Furthermore, the performance of MIQP depends on the geometry of the problem and its execution time can change by orders of magnitude at every control iteration, as will be further explained in section 6.3.


Figure 7: Figure from [16] that represents the exponential growth of the possible modes tree for a single point finger pusher-slider system. The root is a dummy node and, for every other tree level i, each node represents a possible hybrid state for the $i_t h$ state transition. Each possible mode can be represented by a path from the root to one of the tree leaves

Using fast iterations of MPC problems to control a hybrid system is a technique used in other robotics fields, as locomotion, and being able to solve them at a frequency high enough to allow direct implementation into a feedback controller design is currently a problem of interest for the robotics community. A proof of this is that in 2016's Workshop on the Algorithmic Foundations of Robotics (WAFR), held in San Francisco last December, Tedrake presented it as one of the current algorithmic open problems in robotics control problems [28].

3.4 Family of Modes

Responding to the need of obtaining fast solutions to MPC-based control problems Hogan and Rodriguez [16] proposed an heuristic based procedure to approximately solve them called Family of Modes (FOM). In their work, they assumed a finite horizon MPC problem as the one described in the previous section. At every control step the heuristic method explores the optimization problem over only a small set of all the possible combinations of hybrid states to decide the instantaneous controller policy.

The approximate MPC is characterized by a family of modes, which is defined as the set of modes to consider. If the MPC family contains K modes, the MPC will only solve the K optimization sub-problems resulting when enforcing each mode's hybrid state schedule. The optimal solution is approximated by the best solution between the K considered. This solution will consist of the values of $\bar{\mathbf{x}}_i$ and $\bar{\mathbf{u}}_i$ for the optimal control trajectory. The controller will then keep executing the first control policy $\mathbf{u}_1$ and recomputing the approximate MPC problem frequently to take into account new information as the finite horizon advances in time as well as to tackle external interference with the system and errors in the dynamic model.

In his study, Hogan applied FOM to a system consisting on a single finger point pusher and a square slider and he showed that his method was capable of properly tracking straight line trajectories (and composition of straight lines) under external perturbations by using only three modes. The three modes were selected by hand using intuition on the problem structure. They consisted on exploring the three possible hybrid states (using the notation of section 2.3 they were Sticking, Sliding Right, and Sliding Left) on the first step of the MPC and constraining the system to be in a Sticking state for the rest of the steps.

In this simple case it is possible to derive a good set of modes intuitively, but for more complex and higher dimensional problems, this may not be possible. Furthermore, adding more modes to consider into this heuristic doesn't generally translate into better results and, in some cases, it may lead to results far worse than the original one. For some cases, adding new modes to a family that was able to track a line providing asymptotic convergence could even translate into losing this convergence.

We believe the FOM approach shows promising results and may be able to partially address the current needs of manipulation research to solve the MPC based hybrid control problem. At the same time, we consider necessary to show its generality and robustness as well as address some of its limitations for it to be adopted in the robotics community.

In the next section we introduce key concepts on sequential properties of heuristic based methods, that will be useful to describe some of the FOM limitations described above. In chapter 4 we will revisit the main objectives of our study using the concepts presented in previous sections.

3.5 Sequential Properties

On the late 90's, Bertsekas et al. presented a work on the sequential properties of heuristic algorithms [5]. Two important definitions introduced in their study are sequential consistency and sequential improvement. They also introduce the concept of a roll-out algorithm, which holds a close resemblance to the FOM approach to solve MPC problems. At the end of the section we will further discuss this, but we first need to introduce the main definitions and conclusions of the sequential properties study. We will use the same introductory example used in their study because it is similar to the problem we try to solve. The main differences will be mentioned at the end of the section.

Let us consider a graph search problem on a directed graph with node set N and arc set A. For simplicity, an initial node s is defined. For consistency, we will keep the formulation presented in Bertsekas' study, that groups all the search costs in a subset of terminal (with no outgoing arcs) nodes $\bar{N}$ called destinations.

In this formulation, each path from the origin to a destination node $(s, i_1, i_2, ..., i_n, \bar{i})$ (see remark 3.7) has an associated cost $g(\bar{i})$ solely defined by the destination node. We want to find a directed path that starts at the origin s, ends at one of the destination nodes $i \in \bar{N}$, and is such that the cost g(i) is minimized.

Remark 3.7. In order for a sequence $(i_1, i_2, ..., i_n)$ to uniquely define the path created by the composition of the arcs $(i_1, i_2), (i_2, i_3), ..., (i_{n-1}, i_n)$ we need to assume that given an ordered pair of nodes (i, j), there is at most one arc (i, j) with start node i and end node j. This assumption can always be safely made.

Remark 3.8. We allow the node and arc sets, N and A, to contain an infinite number of elements. We require, however, that the number of destination nodes be finite.

The search problem can then be formulated as a DP problem with the nodes corresponding to the states of the DP and the arcs corresponding to the possible state transitions, and the destination nodes corresponding to the terminal states. An heuristic H is then defined as a path construction algorithm that construct a path from a non-destination node to a destination one (remark 3.9) and $H(i)$ is defined for any node $i \in N$ as the cost of the destination node of the path constructed by H:

$$H(i) = g(\bar{i})$$

Remark 3.9. Implicit in this assumption is that for every non-destination node, there exists at least one path starting at that node and ending at some destination node.

Having introduced this definition of heuristic in this search problem, the concept of roll-out algorithm is introduced. The formal definition of roll-out algorithm escapes the scope of this study, but an intuitive definition would be:

Definition 3.10. Roll-out Algorithm:

Given a starting node s and a heuristic H, we define the roll-out algorithm based on H (called *RH*) as following:

1. We consider all downstream neighbours i (with $|i| = n$) of s and we apply H at each possible i.

2. We consider the n paths $(i, i_1, ..., \bar{i})$ obtained from the previous step, each of which has an associated cost $H(i)$.

3. We compute the node i^* which minimizes $H(i)$ and we move to the neighbour i_1^*.

4. The algorithm is then iteratively developed by considering i_1^* the new start node s and repeating the procedure from step 1 until s is a terminal node.

The nature of the roll-out algorithm and the FOM approach are similar in the sense that use simpler base heuristics that would not give satisfying solutions by themselves and keep applying them at every new step, exploring new paths that were not considered when the heuristics were applied to the root node. This study provided some interesting properties that related the heuristic H with its roll-out *RH*. Most of them were properties that guaranteed that *RH* would eventually stop (a problem we don't have in our MPC approach, because $|N|$ and $|A|$ are finite) or variations of the roll-out technique that we will discuss later. In order to properly explain them we need to first introduce two main concepts that helped us to understand some of the limitations of the current FOM formulation and that will be discussed on section 6.4.

Definition 3.11. Sequential Consistency:

The heuristic H is said to be sequentially consistent if for every node i, whenever H generates the path $(i, i_1, ..., i_m, \bar{i})$ starting at i, it also generates the path $(i_1, ..., i_m, \bar{i})$ starting at the node i_1.

Definition 3.12. Sequential Improvement:

Suppose that algorithm H generates, starting at each node $i \in N$, a path $(i, i_1, ..., i_m, \bar{i})$ with the property $H(i) \geq H(i_1)$. Then the heuristic H is said to be sequentially improving.


Figure 8: Excerpt from Bertsekas' study. It shows the graph associated with a one dimensional walk problem. The highlighted path corresponds to the result obtained by applying both the roll-out algorithm RH_1 from example 3.13. A plot of the cost function g(i) is shown on the lower part of the picture.

We are going to borrow two of Bertsekas' examples (associated with Figure 8) to illustrate this two concepts.

Example 3.13. One dimensional walk with a sequentially consistent heuristic:

Consider a person who walks on a straight line in a temporal and spatial discrete world. During each time step, it can move a fixed length unit to the right or to the left (it cannot stay still). At time 0 this person is at position 0 and it walks for a fixed number of time steps N before ending up in a position $\bar{i} \in [-N, N]$. Each terminal point has an associated cost $g(\bar{i})$ and the objective is to minimize the cost of the terminal position. Figure 8 shows a graph representation of the analogue graph search problem, where each node is defined by a state (k, m) where k is the number of steps taken so far and m is the current position. Each node (k, m) has two outgoing arcs: to $(k+1, m-1)$ for the left step and to $(k+1, m+1)$ for the right one.

Let H_1 be an heuristic defined as follows: If H_1 applied at node (k, m), it takes $N - k$ successive steps to the right, terminating at the node $(N, m + N - k)$. Note that H_1 clearly is sequentially consistent.

The algorithm RH_1 applied at node (k, m) uses the roll-out algorithm described above and obtains the path highlighted in Figure 8. In Bertsekas' study it is further proven that the obtained result is always a local minimum.

Example 3.14. Consider the set-up of example 3.13, and let H_2 be defined as follows: when applied at node (k, m), H_2 compares the cost $g(m + N - k)$ (corresponding to taking all of the remaining $N - k$ steps to the right) and the cost $g(m - N + k)$ (corresponding to taking all of the remaining $N - k$ steps to the left). If the first is lower, it takes one step to the right. Otherwise, it takes one step to the left. In Bertsekas' work it is mentioned that H_2 has no sequential consistency propriety, but has instead sequential improvement. It is also shown that RH_2 obtains the global minimum.

Using these definitions, Bertsekas et al. derive some relations between the base heuristics and their associate roll-out procedure. In section 6.4 we propose a method to guarantee the sequential improvement of the FOM method when used with a full horizon instead of a finite one. Our method was inspired in a modification of the traditional roll-out technique proposed by Bertsekas called optimal roll-out.

Consider that at a given roll-out step, the heuristic H generates the optimal path $(i_0, i_1, ..., i_m, \bar{i})$. Then the fortified roll-out algorithm stores the maximal suffix $S_0 = (i_1, ..., i_m, \bar{i})$ and, in the next roll-out iteration, it will consider the paths generated by H together with the path defined by S_0. If the optimal path given by H is not better than the one provided by S_0, the algorithm executes the next step of S_0 and stores the next maximal suffix $S_1 = (i_2, ..., i_m, \bar{i})$. If the path provided by H is better, its first step is executed and the algorithm stores its maximal suffix. This procedure is followed until reaching one of the terminal nodes.

The main idea behind this modified roll-out technique is to simply keep comparing all the possible paths at a certain step with the best path from the previous step. This way the sequential improvement of the algorithm is clearly guaranteed.

4. Goal Revisited

To solve MPC problems fast enough to implement them online in a control problem is an open problem of interest for the robotics manipulation community, as explained in section 3.3. We consider that Hogan's work in the FOM technique is a good way to provide an heuristic-based approximate solution to this problem. The main objective of this study is to bring this method close to the research community by:

1. **Improving the properties of FOM:** In section 6.2 we will discuss the differences in the formulation used in Hogan's study and the one used in this one, as well as why is it fundamental for the proper execution of the FOM method. In section 6.4 we will discuss some of its sequential weaknesses and how we addressed them, and in section 7.3 we will present a new FOM-based method to track complex trajectories.

2. **Comparing it to other techniques:** In section 6.3 we will discuss some of the difficulties we found while trying to compare it to MIQP or MIQP approximations.

3. **Providing a systematic set-up:** In chapter 7 we will present a systematic protocol with an stochastic approach to choose the modes of the family.

4. **Showing the system scalability and robustness:** We want to track more complex trajectories, more complex systems, and use models of a simple system to solve a problem that uses a more complex

one. Some of these goals will be reached in this study and will be discussed through chapters 6 and 7. We will also present our current line of research and experiments we are currently developing for ISRR in chapter 9.

In order to achieve this goals, we provided quantitative tools to evaluate the obtained trajectories that will be presented in section 6.1 and we implemented a general simulation interface (presented in chapter 6) that allowed us to easily formulate and solve general trajectories and hybrid systems.

The formulation used in our study will be presented in the next chapter, together with all its required theoretic concepts.

5. Pusher-Slider Formulation

In this section we are going to show all the necessary steps to formulate the FOM problem for the case of a rectangular pusher slider-system with a single finger point pusher.

5.1 Nomenclature Summary

Property	Symbol	Value
Coeff. friction (Pusher-Slider)	μ_p	0.3
Coeff. friction (Slider-Ground)	μ_g	0.35
Mass of the slider (Kg)	m	1.05
Length of the slider (m)	a	0.09
Width of the slider (m)	b	0.09

Table 1: Properties of the slider and the support surface

Property	Symbol
Length of the finite horizon	N
Time step	h
State cost matrix	$\mathbf{Q}$
Control input cost matrix	$\mathbf{R}$
Final state cost matrix	$\mathbf{Q}_f$

Table 2: MPC variable summary

Property	Symbol
Position of the slider's CM	$(x, y)^T$
Slider orientation	θ
Slider inertial frame	F
Pusher position in F	$(p_x, p_y)^T$
Pusher velocity in F	$(v_x, v_y)^T$
Trajectory state	$\mathbf{x} = (x, y, \theta, p_y)^T$
Control input	$\mathbf{u} = (v_x, v_y)^T$

Table 3: Pusher-Slider system variable summary

To derive the motion equations and define the MPC state and controller variables, we are going to use the notation used in [16]. We will first define an inertial reference frame F centred on the slider center of mass that will move and with it. The pose $\mathbf{q}$ of the slider's center of mass will be defined in a fixed world frame $\mathbf{q} = (x, y, \theta)^T$ and the pose $\mathbf{p}$ of the pusher will be defined relative to the slider center of mass $\mathbf{p} = (p_x, p_y)^T$ as shown in Figure 9. The state variable used to define a trajectory is $\mathbf{x} = (x, y, \theta, p_y)^T$ and the control input is the velocity of the pusher defined in the center of mass frame $\mathbf{u} = \mathbf{v}_p = (v_x, v_y)^T$.


Figure 9: Figure from [16] that shows the inertial frame used to define some of the state variables

When talking about hybrid states and modes, we will usually use the following naming convention:

- *S*: Sticking State.

- *R*: Sliding Right State.

- *L*: Sliding Left State.

- $\mathbf{S}^N$: Mode defined by N S states $(S, \overset{N}{\cdots}, S)$.

- $\mathbf{R}_i^N$: Mode defined by $N - 1$ states S and one state R in position i. For example, $\mathbf{R}_1^3 = (R, S, S)$.

- $\mathbf{L}_i^N$: Analogue to $\mathbf{R}_i^N$. Mode defined by $N - 1$ states S and one state L in position i.

5.2 Quasistatic Pushing Formulation

Quasistatic analysis consists in assuming balance among contact wrenches, gravitation and other applied forces, while neglecting inertial wrenches. This analysis is often taken in manipulation problems, where the trajectory tracking speed changes don't incur into high inertial forces.

The slider of our pusher-slider system with a single finger point pusher is affected by two wrenches: the wrench the pusher applies on the slider ($\mathbf{f}_p$) and the wrench the support surface applies to it ($\mathbf{f}_g$). Under the quasistatic assumption $\mathbf{f}_p = -\mathbf{f}_g$.

This assumption will reduce the dimension of the state variables in the motion equations derived in section 5.5 and will allow as to ignore considering a separation hybrid state in our hybrid model of the pusher slider system (if inertial forces were non-negligible the slider could separate away from the pusher while tracking complex trajectories with points of high curvatures). Furthermore, combining the quasistatic assumption with the limit surface of the slider-ground contact, we can derive velocity based motion equations and friction constraints, as will be shown in section 5.5. The robot arms we had available for this study are velocity-controlled, so this formulation simplifies the experimental implementation of the controller.

5.3 Ellipsoidal Approximation Of The Limit Surface

Given a pressure distribution, the associated surface distribution can be numerically computed. Unluckily, the pressure distribution is usually not known exactly and changes over time. For our this study we will assume that the slider and the supporting surface are completely flat, providing a uniform pressure distribution, so $f_n = \dfrac{gm}{A}$, where f_n is the magnitude of the normal support force the ground exerts on the slider and m and A are the slider mass and area, respectively. Furthermore, we will use an approximation to the actual limit surface proposed by Lee and Cutkosky on 1991 [18] based on a study made on 1988 by Howe, Kao, and Cutkosky that showed the suitability of using an elliptic approximation for dry Coulomb friction [17]. We will add further simplifications to their ellipsoid derivation. For a single contact point, their derivation can be simplified, as the limit surface will already be oriented in the ellipsoid major axis. Assuming uniform pressure distribution also simplifies their computation of the maximum momentum around the instantaneous center of rotation. The final procedure to obtain the LS ellipsoid is as follows:

1. Compute the maximum momentum value. This occurs when the object is in pure rotation, so we can use equation 4 considering $\mathbf{r}_{IC}$ to be the center of mass (CM) of the slider:

$$n_{max} = \left| -\mu \, sign(\dot{\theta}) \int_F \mathbf{r} \, \frac{\mathbf{r} - \mathbf{r}_{CM}}{|\mathbf{r} - \mathbf{r}_{CM}|} p(\mathbf{r}) dA \right| = \mu \frac{gm}{A} \int_F ||\bar{\mathbf{r}}|| dA \tag{7}$$

where $\bar{\mathbf{r}}$ is the vector from $\mathbf{r}_{CM}$ to $\mathbf{r}$. This gives us the value of the moment semi-axis (the maximum and minimum values of m in the ellipse are going to be $\pm n_{max}$).

2. Similarly, to obtain the force semi-axes, we need the maximum force in the x and y directions, which is going to be $f_{max} = f_n$.

3. The final expression for the ellipsoid is:

$$\left(\frac{f_x}{f_{max}}\right)^2 + \left(\frac{f_y}{f_{max}}\right)^2 + \left(\frac{n}{n_{max}}\right)^2 = 1 \tag{8}$$


Figure 10: Figure from [18] that shows the experimental limit surface for different contact assumptions (that result in different pressure distributions). For the uniform pressure distribution the ellipsoidal approximation describes accurately the LS

5.4 Motion Cone

In Hogan's study, a velocity controlled pusher was used. For this reason the motion equations for the MPC problem were formulated in the velocity space. For the case of a single point finger pushing, we can obtain an expression for the friction cone in the velocity space called motion cone (MC) [22] that will be useful to formulate the motion equations and the friction constraints.

As opposed to friction cone, the motion cone is not symmetric with respect to the normal of the contact surface at the point of contact. The cone can be determined by using two angles $\gamma_t(\mathbf{p})$ and $\gamma_b(\mathbf{p})$. Using

the ellipsoidal approximation of the LS, the expression of the motion cone for flat sliders can be given as:

$$\gamma_t = \frac{\mu_p c^2 - p_x p_y + \mu_p p_x^2}{c^2 + p_y^2 - \mu_p p_x p_y} \tag{9}$$

$$\gamma_b = \frac{-\mu_p c^2 - p_x p_y - \mu_p p_x^2}{c^2 + p_y^2 + \mu_p p_x p_y} \tag{10}$$

where p_x and p_y have been defined in section 5.1 and $c = \frac{f_{max}}{n_{max}}$.


Figure 11: Figure from [16] with a representation of the motion cone

We will define $\gamma := \frac{v_y}{v_x}$. If the velocity vector lies inside the motion cone, the pusher will stick. If $\gamma > \gamma_t$ the slider will slip to the right with respect to the pusher. Finally, if $\gamma < \gamma_b$ the slider will slip to the left with respect to the pusher. In summary the friction constraints can be expressed for each hybrid state as:

$$\textbf{Sticking:} \quad \textbf{u} \in MC : \begin{cases} v_y \leq \gamma_t v_x \\ v_y \geq \gamma_b v_x \end{cases} \tag{11}$$

$$\textbf{Sliding Right:} \quad \textbf{u} > MC : v_y > \gamma_t v_x \tag{12}$$

$$\textbf{Sliding Left:} \quad \textbf{u} < MC : v_y < \gamma_b v_x \tag{13}$$

In our single finger point pusher system the three possible hybrid states of the system will be these. The quasistatic assumption guarantees that the slider doesn't move away from the pusher by virtual inertial forces and we will impose constraints on the pusher to prevent it to move away from the slider: the normal velocity of the pusher with respect to the contact surface will be non-negative and the tangential

displacement from the contact surface center (p_y) will be constrained to prevent the pusher from slipping through the slider sides.

5.5 Motion Equations

The motion equations are adapted from the ones used in Hogan's work. They were originally inspired in a pusher-slider system study by Lynch et al. [19] and can be expressed as follows:

$$\dot{\mathbf{x}} = \begin{cases} \mathbf{f}_1(\mathbf{x}, \mathbf{u}) \text{ if } \mathbf{u} \in \text{MC} \\ \mathbf{f}_2(\mathbf{x}, \mathbf{u}) \text{ if } \mathbf{u} > \text{MC} \\ \mathbf{f}_3(\mathbf{x}, \mathbf{u}) \text{ if } \mathbf{u} < \text{MC} \end{cases} \tag{14}$$

$$\mathbf{f}_i(\mathbf{x}, \mathbf{u}) = \begin{pmatrix} \mathbf{CQP}_i \\ \mathbf{b}_i \\ \mathbf{c}_i \end{pmatrix} \mathbf{u} := \hat{\mathbf{B}}_j(\mathbf{x}) \, \mathbf{u} \tag{15}$$

where $\mathbf{C} = \begin{pmatrix} cos(\theta) & -sin(\theta) \\ sin(\theta) & cos(\theta) \end{pmatrix}$ and $\mathbf{Q} = \dfrac{1}{c^2 + p_x^2 + p_y^2} \begin{pmatrix} c^2 + p_x^2 & p_x p_y \\ p_x p_y & c^2 + p_y^2 \end{pmatrix}$.

$\mathbf{P}_i$, $\mathbf{b}_i$ and $\mathbf{c}_i$ are mode-dependant matrices and take the following values:

$\mathbf{b}_1 = \dfrac{1}{c^2 + p_x^2 + p_y^2} \begin{pmatrix} -p_y & p_x \end{pmatrix}$, $\mathbf{b}_2 = \dfrac{1}{c^2 + p_x^2 + p_y^2} \begin{pmatrix} -p_y + \gamma_t p_x & 0 \end{pmatrix}$, $\mathbf{b}_3 = \dfrac{1}{c^2 + p_x^2 + p_y^2} \begin{pmatrix} -p_y + \gamma_b p_x & 0 \end{pmatrix}$

$\mathbf{c}_1 = \begin{pmatrix} 0 & 0 \end{pmatrix}$, $\qquad \mathbf{c}_2 = \begin{pmatrix} -\gamma_t & 1 \end{pmatrix}$, $\qquad \mathbf{c}_3 = \begin{pmatrix} -\gamma_b & 1 \end{pmatrix}$

$\mathbf{P}_1 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, $\qquad \mathbf{P}_2 = \begin{pmatrix} 1 & 0 \\ \gamma_t & 0 \end{pmatrix}$, $\qquad \mathbf{P}_3 = \begin{pmatrix} 1 & 0 \\ \gamma_b & 0 \end{pmatrix}$

where i = 1, 2, 3 correspond to the states of Sticking, Sliding Right and Sliding Left respectively. For simplicity, we do not consider a state where the pusher is separated from the slider for the controller modelling. In the controller we impose constraints that forbid the pusher to move outside of the slider side borders and the under the quasistatic assumption, the slider will not separate from the pusher by it inertia.

On the other hand the possibility of separation is included in the simulator, in order to represent the experimental set-up more accurately and to detect possible errors in the controller design.

The motion equations are going to be used to define constraints in the MPC problem and, in turn, the MPC formulation is going to be fed into a commercial optimizer to obtain the optimal values of $\mathbf{x}$ and $\mathbf{u}$ that fulfil all the constraints. The fact that the motion equations are nonlinear would require using a nonlinear optimizer, which would represent too much computational cost to solve the problem online at every control iteration. In order to obtain a more tractable formulation the motion equations will be linearized around a nominal trajectory $\mathbf{x}^*(t)$ and control input $\mathbf{u}^*(t)$. We will use $(.)^*$ to denote elements of the nominal problem and $\bar{(.)}$ to denote perturbations around the nominal value. Using this notation, the linearization of the motion equation 14 is:

$$\dot{\bar{\mathbf{x}}} = \mathbf{A}_i(\mathbf{x}^*, \mathbf{u}^*) \, \bar{\mathbf{x}} + \mathbf{B}_i(\mathbf{x}^*, \mathbf{u}^*) \, \bar{\mathbf{u}} + \mathbf{f}_i(\mathbf{x}^*, \mathbf{u}^*) - \dot{\mathbf{x}}^* \tag{16}$$

$\mathbf{A}_i(\mathbf{x}', \mathbf{u}') = \dfrac{\partial \mathbf{f}_i(\mathbf{x}, \mathbf{u})}{\partial \mathbf{x}} \bigg|_{\mathbf{x}', \mathbf{u}'}$, $\mathbf{B}_i(\mathbf{x}', \mathbf{u}') = \dfrac{\partial \mathbf{f}_i(\mathbf{x}, \mathbf{u})}{\partial \mathbf{u}} \bigg|_{\mathbf{x}', \mathbf{u}'} = \hat{\mathbf{B}}_i(\mathbf{x}')$,with $\bar{\mathbf{x}} = \mathbf{x} - \mathbf{x}^*$ and $\bar{\mathbf{u}} = \mathbf{u} - \mathbf{u}^*$

It should be noted that equation 16 is different from the linearized motion equation presented in Hogan's work. This fact will be discussed further in section 6.2 but the main idea is that the term $\mathbf{f}_i(\mathbf{x}^*, \mathbf{u}^*) - \dot{\mathbf{x}}^*$ cannot generally be cancelled as in other traditional control problems because $\mathbf{f}_i(\mathbf{x}^*, \mathbf{u}^*)$ will not describe the instantaneous motion of the nominal trajectory if the nominal trajectory is in a hybrid state different from the one defined by i.

Similarly to the linearization of the motion equation, the friction constraint equations from 11 to 13 can also be linearized with respect to $\bar{\mathbf{x}}$. For consistency with the variables used for the motion equations ($\bar{\mathbf{x}}$ and $\bar{\mathbf{u}}$) v_x and v_y are expressed as $v_x^* + \bar{v}_x$ and $v_y^* + \bar{v}_y$ respectively:

$$\textbf{Sticking:} \quad \begin{cases} (v_y^* + \bar{v}_y) \leq (\gamma_t(\mathbf{x}^*) + \mathbf{C}_t(\mathbf{x}^*)\,\bar{\mathbf{x}})\,(v_x^* + \bar{v}_x) \\ (v_y^* + \bar{v}_y) \geq (\gamma_b(\mathbf{x}^*) + \mathbf{C}_b(\mathbf{x}^*)\,\bar{\mathbf{x}})\,(v_x^* + \bar{v}_x) \end{cases} \tag{17}$$

$$\textbf{Sliding Right:} \quad (v_y^* + \bar{v}_y) > (\gamma_t(\mathbf{x}^*) + \mathbf{C}_t(\mathbf{x}^*)\,\bar{\mathbf{x}})\,(v_x^* + \bar{v}_x) \tag{18}$$

$$\textbf{Sliding Left:} \quad (v_y^* + \bar{v}_y) < (\gamma_b(\mathbf{x}^*) + \mathbf{C}_b(\mathbf{x}^*)\,\bar{\mathbf{x}})\,(v_x^* + \bar{v}_x) \tag{19}$$

where $\mathbf{C}_t(\mathbf{x'}) = \dfrac{\partial \gamma_t}{\partial \bar{\mathbf{x}}}\bigg|_{\mathbf{x'}}$ and $\mathbf{C}_b(\mathbf{x'}) = \dfrac{\partial \gamma_b}{\partial \bar{\mathbf{x}}}\bigg|_{\mathbf{x'}}$

This notation allows us to write all the linearized motion cone constraints in a condensed form suitable for a linear programming optimizer, as equation 16:

$$\mathbf{E}_i(\mathbf{x}^*, \mathbf{u}^*)\,\bar{\mathbf{x}} + \mathbf{D}_i(\mathbf{x}^*)\,\bar{\mathbf{u}} \leq g_i(\mathbf{x}^*, \mathbf{u}^*) \tag{20}$$

$$\mathbf{E}_1(\mathbf{x'}, \mathbf{u'}) = v_y' \begin{pmatrix} -\mathbf{C}_t(\mathbf{x'}) \\ \mathbf{C}_b(\mathbf{x'}) \end{pmatrix}, \qquad \mathbf{E}_2(\mathbf{x'}, \mathbf{u'}) = v_y'\mathbf{C}_t(\mathbf{x'}), \qquad \mathbf{E}_3(\mathbf{x'}, \mathbf{u'}) = -v_y'\mathbf{C}_b(\mathbf{x'})$$

$$\mathbf{D}_1(\mathbf{x'}) = \begin{pmatrix} -\gamma_t(\mathbf{x'}) & 1 \\ \gamma_b(\mathbf{x'}) & -1 \end{pmatrix}, \qquad \mathbf{D}_2(\mathbf{x'}) = \begin{pmatrix} -\gamma_t(\mathbf{x'}) & -1 \end{pmatrix}, \qquad \mathbf{D}_3(\mathbf{x'}) = \begin{pmatrix} -\gamma_b(\mathbf{x'}) & 1 \end{pmatrix}$$

$$\mathbf{g}_1(\mathbf{x'}, \mathbf{u'}) = \begin{pmatrix} -v_y' + \gamma_t(\mathbf{x'})v_x' \\ v_y' - \gamma_b(\mathbf{x'})v_x' \end{pmatrix}, \qquad \mathbf{g}_2(\mathbf{x'}, \mathbf{u'}) = v_y' - \gamma_t(\mathbf{x'})v_x' - \epsilon, \qquad \mathbf{g}_3(\mathbf{x'}, \mathbf{u'}) = -v_y' + \gamma_b(\mathbf{x'})v_x' - \epsilon$$

Remark 5.1. ϵ is just a small positive scalar used to express strict inequalities, as the optimizer we used only supported non-strict ones.

5.6 MPC And MIQP Formulation

As explained in section 3.3, each MPC mode fixed sub-problem will compute the optimal values for the state and control variables for a discrete finite horizon of N steps, a time step h, and a given initial position $\bar{\mathbf{x}}_0$ to minimize the cost equation provided in equation 6, subject to motion and cone constraints. We will use the perturbation dynamics, so the state and control variables will be $(\bar{\mathbf{x}}_1, ..., \bar{\mathbf{x}}_N)$ and $(\bar{\mathbf{u}}_0, ..., \bar{\mathbf{u}}_{N-1})$; and the initial state will be $\mathbf{x}_0$. As MPC considers a discrete framework, we need to discretize the motion equations from previous section. After a first order discretization, the MPC constraints at step j are:

$$\text{if } j = 0 : \begin{cases} \bar{\mathbf{x}}_1 = \bar{\mathbf{x}}_0 + h\mathbf{B}_{i0}(\mathbf{x}_0)\,\bar{\mathbf{u}}_0 + h\mathbf{f}_{i0} + \mathbf{x}_0^* - \mathbf{x}_1^* \\ \mathbf{D}_{i0}\,\bar{\mathbf{u}}_0 \leq \mathbf{g}_{i0} \end{cases} \tag{21}$$

$$\text{if } j > 0 : \begin{cases} \bar{\mathbf{x}}_{j+1} = (\mathbf{I} + h\mathbf{A}_{ij})\,\bar{\mathbf{x}}_j + h\mathbf{B}_{ij}\,\bar{\mathbf{u}}_j + h\mathbf{f}_{ij} + \mathbf{x}_j^* - \mathbf{x}_{j+1}^* \\ \mathbf{E}_{ij}\,\bar{\mathbf{x}}_j + \mathbf{D}_{ij}\,\bar{\mathbf{u}}_j \le \mathbf{g}_{ij} \end{cases} \tag{22}$$

where $\mathbf{A}_{ij} = \mathbf{A}_i(\mathbf{x}_j^*, \mathbf{u}_j^*)$, and $\mathbf{E}_{ij} = \mathbf{E}_i(\mathbf{x}_j^*, \mathbf{u}_j^*)$, using the notation from section 5.5. $\mathbf{f}_{ij}$ is defined differently depending on the iteration step:

$$\mathbf{f}_{ij} = \begin{cases} \mathbf{f}_i(\mathbf{x}_j, \mathbf{u}_j^*), & (\text{if } j = 0) \\ \mathbf{f}_i(\mathbf{x}_j^*, \mathbf{u}_j^*), & (\text{if } j > 0) \end{cases}$$

$\mathbf{B}_{ij}$, $\mathbf{D}_{ij}$, $\mathbf{E}_{ij}$, and $\mathbf{g}_{ij}$ are defined in an analogue way, evaluating their section 5.5 counterparts on $\mathbf{x}_j$ for the first step and on $\mathbf{x}_j^*$ for the following ones.

Remark 5.2. For the case j = 0, $\bar{\mathbf{x}}_0$ is known value, so the first order discrete version of equation 14 becomes:

$$\mathbf{x}_1 = \bar{\mathbf{x}}_1 + \mathbf{x}_1^* = h\mathbf{f}_i(\mathbf{x}_0, \mathbf{u}_0) + \mathbf{x}_0 = h\hat{\mathbf{B}}_i(\mathbf{x}_0)(\bar{\mathbf{u}}_0 + \mathbf{u}_0^*) + \bar{\mathbf{x}}_0 + \mathbf{x}_0^* \tag{23}$$

and by isolating $\bar{\mathbf{x}}_1$ we can obtain the linear motion equation in expression 21 without having to apply a linearization step. In a similar way, the friction constraints can also be expressed as linear constraints with respect to $\bar{\mathbf{u}}_0$

Let $m = (s_0, ..., s_N)$ denote the hybrid mode where at step j of the fixed mode MPC sub-problem the system is in the hybrid state s_j. Then, the sub-problem constraints are formulated by substituting i with s_j in equations 21 and 22. Then the optimization problem is fed into the commercial optimizer Gurobi [2] to obtain the optimal values of $(\bar{\mathbf{x}}_1, ..., \bar{\mathbf{x}}_N)$ and $(\bar{\mathbf{u}}_0, ..., \bar{\mathbf{u}}_{N-1})$, together with the optimal cost function. The full MPC problem solution is then approximated to be the best sub-problem solution.

As will be explained in section 6.3, we also implemented an MIQP solver to compare our FOM against. The formulation we have presented in this section has already linear constraints and quadratic cost function, so we can apply quadratic programming solvers to them. Gurobi allowed us to formulate the mode dependent constraints by using binary decision variables and a big-M notation. The resulting formulation for the motion and friction constraints is as follows:

$$\text{if } j = 0 : \begin{cases} \mathbf{x}_1 \le (\bar{\mathbf{x}}_0 + h\mathbf{B}_{i0}(\mathbf{x}_0)\,\bar{\mathbf{u}}_0 + h\mathbf{f}_{i0} + \mathbf{x}_0^* - \mathbf{x}_1^*) + M(1 - z_{i0}) \\ -\mathbf{x}_1 \le -(\bar{\mathbf{x}}_0 + h\mathbf{B}_{i0}(\mathbf{x}_0)\,\bar{\mathbf{u}}_0 + h\mathbf{f}_{i0} + \mathbf{x}_0^* - \mathbf{x}_1^*) + M(1 - z_{i0}) \\ \mathbf{D}_{i0}\bar{\mathbf{u}}_0 \le \mathbf{g}_{i0} + \mathbf{1}_{2\times 1}M(1 - z_{i0}) \end{cases} \tag{24}$$

$$\text{if } j > 0 : \begin{cases} \bar{\mathbf{x}}_{j+1} \le [(\mathbf{I} + h\mathbf{A}_{ij})\,\bar{\mathbf{x}}_j + h\mathbf{B}_{ij}\,\bar{\mathbf{u}}_j + h\mathbf{f}_{ij} + \mathbf{x}_j^* - \mathbf{x}_{j+1}^*] + M(1 - z_{ij}) \\ -\bar{\mathbf{x}}_{j+1} \le -[(\mathbf{I} + h\mathbf{A}_{ij})\,\bar{\mathbf{x}}_j + h\mathbf{B}_{ij}\,\bar{\mathbf{u}}_j + h\mathbf{f}_{ij} + \mathbf{x}_j^* - \mathbf{x}_{j+1}^*] + M(1 - z_{ij}) \\ \mathbf{E}_{ij}\bar{\mathbf{x}}_j + \mathbf{D}_{ij}\bar{\mathbf{u}}_j \le \mathbf{g}_{ij} + \mathbf{1}_{2\times 1}M(1 - z_{ij}) \end{cases} \tag{25}$$

$$z_{1j} + z_{2j} + z_{3j} = 1 \tag{26}$$

The equality constraints $a = b$ (where a and u are functions of $\mathbf{x}$ and $\mathbf{u}$) have been broken into two inequality ones $a \le b$ and $-a \le -b$ and each mode independent constraint $a \le b$ has been modified into $a \le b + M(1 - z_{ij})$ for an appropriate binary variable z_{ij}. If the sub-problem is in state i at step j, $z_{ij} = 1$ and the term $M(1 - z_{ij})$ is zero and the original inequality constrains applies. If the sub-problem is not in that state at step j, $z_{ij} = 0$ and the inequality becomes $a \le b + M$. If the value of M is high enough (specifically, if it is higher than any possible value of the expression a - b) the inequality becomes redundant. Equation 26 enforces that the MPC sub-problems are in one and only one hybrid state at each step.

6. Simulation Interface

The simulations of Hogan's study were mainly designed for an specific choice of modes and a specific trajectory to follow. In this study, we wanted to study general properties of the FOM problem, to generalize the results to more complex pusher-slider systems and trajectories, and to compare our different FOM strategies to solve the MPC problem with other conventional methods as MIQP. In order to satisfy this needs, we decided to implement a general simulation tool. The code main structure can be described as follows:

An abstract state interface is implemented to provide the dynamics and cone constraints of the state as well as their linearized versions at a certain objective state space and control input. To define a certain experimental problem, the dynamics function and constraint equations are derived depending on a general state space variable $\mathbf{x}$ and controller action $\mathbf{u}$ and the necessary state instances are implemented and a mode is then defined by a vector of state handles. An MPC solver class is also implemented to get an initial condition and an objective trajectory and return a control input. Each MPC solver implementation can include its own member states or modes (as well as other optimization set-up variables) and internally decides how to handle them.

At the same time, an Euler integration simulator is implemented. This simulator uses its own state classes to describe the dynamics of the pusher slider system. These states are independent from the optimization ones so that discrepancies between model and reality can be simulated (for example the MPC solver can use linearized dynamics or cone constraints while the Euler integrator uses the non-linearized version, which is more similar to the real experimental set-up than using linearized dynamics in both places). The Euler integration class also holds a member MPC solver that returns the control input to be applied. The integrator then computes the value of the state variable derivative $\dot{\mathbf{x}}$ and uses a simple quadrature to approximate the next state variable value.

The controller (MPC solver) and simulator are decoupled, so they can consider different hybrid states. This allows us to consider more realistic models for the euler integration step of the simulation (we always use non-linearized motion equations for the integration step), to use different parameters in their states (to test controllers designed for a certain slider shape, μ_g, and μ_p together with a simulator set-up with different parameters, to show the robustness of the controller to uncertainties in the models).

6.1 Local, Global and MPC Costs

During the previous work at MCubelab, there was no systematic and objective way to evaluate the global trajectory resulting from the FOM control law, the exit or failure of an experiment was mainly determined in a subjective binary way. If the pusher slider didnt show appreciable divergence (at the naked eye) from the planned trajectory it was considered a success, otherwise it was a failure. In a similar fashion if the system was perturbed (by an external interaction), the success lied in recovering the original trajectory in a reasonable time determined subjectively. The main goal of the study was just to demonstrate that FOM could effectively control the pusher slider system to track previously planned trajectories, and that all this process could be done online. Because of this proof-of-concept nature, this approach was sufficient.

For this thesis we wanted a more objective way to evaluate how good a control trajectory was and this required a certain notion of metric or score for the global trajectory. We tried to keep coherence between the global and local cost functions, but removed the $\bar{\mathbf{x}}_N^T \mathbf{Q}_f \bar{\mathbf{x}}_N$ term of the cost function. As presented in section 3.3, this term is used to account for the information missing because of the finite horizon. When

scoring the final trajectory, we have the full horizon information, so this term is no longer required. The obtained cost function is $S = \sum_{i=0}^{M} (\bar{\mathbf{x}}_i^T \mathbf{Q} \bar{\mathbf{x}}_i + \bar{\mathbf{u}}_i^T \mathbf{R} \bar{\mathbf{u}}_i)$, where M is the number of euler integration steps. We will denote by local cost over time each of the $\bar{\mathbf{x}}_i^T \mathbf{Q} \bar{\mathbf{x}}_i + \bar{\mathbf{u}}_i^T \mathbf{R} \bar{\mathbf{u}}_i$ values.

It is important to note that even if this global cost function is used score how close the obtained trajectory was from the desired one, it is not a direct indicator of the MPC solving performance. We will denote the cost function provided in section 3.3 as MPC cost. Then optimizing the MPC cost at each iteration could incur in a larger optimal cost for the global trajectory. If we are currently at the state x_i and we consider two modes M_1 and M_2 with local costs c_1 and c_2, with $c_1 < c_2$ there's no guarantee that executing the first step of the trajectory obtained from mode M_1 will translate in a smaller global cost function. Some causes for this effect are exposed and fixed in 6.4 but, even after fixing them and considering an ideal world scenario (where the models are completely accurate and there's no sensor noise), the fact that the receding finite horizon adds new information to the problem at every control iteration implies that the optimal state and control pair $(\mathbf{x}_{i+1}, \mathbf{u}_i)$ computed at the control iteration i could technically lose their optimal property when the information from control iteration $i + 1$ is introduced.

This effect is common in all receding finite horizon problems and there's usually no way to completely solve it. In order to evaluate the MPC solving performance of our method, some of the simulation results we expose in the following sections show a plot of the local MPC cost over time in addition to a graphical representation of the trajectory and its global or local cost. It will be shown that, when extending a family F_1 into a family F_2 by adding more modes, the local cost of F_2 descends with respect ot F_1 in the first bifurcation of the local MPC cost over time plot. This will be further discussed in section 6.5.

6.2 Nominal Mode

In section 5.5 we noted that, as opposed to the previous study, we do not simplify the term $\mathbf{f}_i(\mathbf{x}^*, \mathbf{u}^*) + \dot{\mathbf{x}}^*$ in equation 16. $\mathbf{x}^*$ can be described by the same motion equations as $\mathbf{x}$ using the hybrid state friction constraints to determine the hybrid state, but it is important to realize that, once we constrain out system to be in a state i at a time instant t_j, the specific motion equation that describes this instant-state system $\mathbf{f}_i(\mathbf{x}_j, \mathbf{u}_j)$ may not describe the instantaneous motion of the nominal trajectory. The nominal trajectory will clearly be described by $\mathbf{f}_{i'}(\mathbf{x}_j^*, \mathbf{u}_j^*)$, but there's no guarantee that $i = i'$. The error introduced by assuming $\mathbf{f}_i(\mathbf{x}^*, \mathbf{u}^*) + \dot{\mathbf{x}}^* = 0$ can be critical for the performance of the controller. In Figure 12 we can see the behaviour of the FOM approach when trying to track a circle using a family of modes learned using the procedure described in section 7.2. The red slider follows the nominal trajectory. The trajectory followed by its center of mass trajectory is represented with the black line. The blue slider follows the simulated trajectory and its center of mass follows the blue line. The last position of the simulated slider is highlighted and the other snapshots are faded away. The considered optimal paths the MPC considers are shown at the center of each simulated slider. The chosen path is the yellow one. As can be seen the obtained trajectory diverges from the nominal one. Furthermore, it can be seen how the optimal finite horizon trajectory computed in the MPC problem is completely different from the realized trajectory. In Figure 13 the correspondent local cost over time is shown.

The possibility of the hybrid states associated the nominal trajectory to be different from the ones associated with the real trajectory has another important implication. The perturbation state motion equation 16 has the nominal trajectory as an equilibrium point only when $i = i'$. Otherwise, if we consider $\mathbf{x} = \mathbf{x}^*$ and $\mathbf{u} = \mathbf{u}^*$, we obtain $\bar{\mathbf{x}} = \mathbf{A}_i \cdot \mathbf{0} + \mathbf{B}_i \cdot \mathbf{0} + \mathbf{f}_i(\mathbf{x}^*, \mathbf{u}^*) - \dot{\mathbf{x}}^* = \mathbf{f}_i(\mathbf{x}^*, \mathbf{u}^*) - \mathbf{f}_{i'}(\mathbf{x}^*, \mathbf{u}^*)$. If none of the MPC modes is the one that the nominal trajectory follows, the solution $\bar{\mathbf{x}}_j = \bar{\mathbf{u}}_j = \mathbf{0}$ will never be feasible,

Figure 12: Snapshot from the video simulating the trajectory obtained when tracking a circular trajectory provided by equation , with $d = 0.0249$. with the 8 best modes obtained using the procedure presented in section 7.2.

and we will never be able to stabilize the planned trajectory on the nominal one, even if the initial state is on the nominal trajectory ($\mathbf{x}_0 = \mathbf{0}$). In order to force the MPC to consider the nominal mode, we can simply compute it at every MPC step (it can be easily determined applying the friction constraints to each pair $\mathbf{x}_j^*$, $\mathbf{u}_j^*$) and add it as one of the modes of the family.

6.3 MIQP and Clustered MIQP

We wanted to compare the solutions provided by FOM with traditional MPC solving techniques. In order to do that we decided to implement the MIQP technique using the commercial solver Gurobi. This solver was already used in Hogan's work in a force-based formulation prior to the final velocity-based one. This formulation was not presented in the final paper, but it allowed to formulate an MIQP problem that was able to run in an average of 3Hz. It should be mentioned that the execution time at every control iteration could range between few milliseconds to 10 seconds depending on the initial condition of the MPC. This is not a problem for simulation, but in a real control environment the frequency at which you can update the controller is not bounded by the average frequency at which you can solve the MPC, but by the worst case scenario one. Otherwise the controller may try to update the control input before it is computed by

Figure 13: Local cost over time associated to the trajectory in Figure 12

the optimizer. There may be solutions to this effect in other formulations, but for the one used in Hogan's work the MIQP controller could run at more than 0.1Hz.

We were expecting a similar behaviour in the velocity-based formulation, but it took several orders of magnitude more time to solve each iteration. In order to simulate 500 simulation steps with a finite horizon of 7 steps it took an average of 6 hours of execution time, which implies that the maximum control frequency we could apply would probably not surpass 0.02Hz. We were expecting to have an "ideal" MIQP result to compare FOM with but, as will be shown in section 6.5, the results obtained with a small finite horizon for the MIQP did not track properly many of the non-trivial nominal trajectories and initial perturbations we fed into it.

We looked at techniques to speed up the MIQP execution in the Gurobi website and we found out that our current big-M formulation worked better (both in execution time and in quality of the obtained trajectory) as smaller the M parameter was. As explained in 5.6 each M parameter used in a scalar inequality $Ax \leq b + M(1 - z)$ must be greater than $\max_{x \in D}(Ax - b)$ (where D is the domain of possible values of the x) so that the inequality is rendered redundant when $z = 0$. We implemented a version of MIQP that given a vectorial constraint as the ones used in Hogan's formulation $A\vec{x} \leq \vec{b} + M(1 - z) \cdot (1, ..., 1)^T$ replaces the arbitrary constant M by a vector $\vec{M}$ where each component's value is $M_i = \max_{x_i \in D}(A_i x_i - b_i)$. We were not able to improve the execution time, but we were able to obtain better results for the MIQP. The results showed in the next section for the MIQP will use this optimized big-M formulation.

In order to compare FOM to other methods using a longer time horizon, we implemented another MPC solving method based on MIQP that we called Clustered MIQP. Given a standard MPC formulation, this method works in the following way:

- Consider a clustering factor c, a finite horizon $N \cdot c$ and a set of h hybrid states $\{H_1, ..., H_h\}$.

- The finite horizon is then divided into N clusters of c steps: $(C_1, ..., C_N)$.

- The state space is extended in a similar way as in section 5.6 with 3 binary variables per cluster for a total of $3N$ variables.

- Each variable z_{1i}, z_{2i}, and z_{3i} is then associated with the hybrid state of the cluster C_i (sticking, sliding down and sliding up respectively).

- The constraint $z_{1i} + z_{2i} + z_{3i} = 1$ is enforced for all i.

- The problem constraints are formulated with big M notation as in section 5.6 and the problem is fed into a Gurobi MIQP solver.

As we exposed in section 3.3, the MIQP formulation doesn't escalate well with the number of binary variables added and clustered MIQP allows us to reduce the binary variable dimensionality of the optimization problem by keeping the same state during all the steps of the same cluster. The execution time of this new problem is similar to the execution time of a MIQP problem of N steps but the controller inputs can be different at each of the Nc steps, providing a more flexible trajectory than the one given by the MIQP problem of N steps.

We intended to use this method to approximate MIQP problems with long horizons by using few clusters and we expected to obtain a better trajectory than the ones obtained with FOM. As will be discussed later, the results of this MIQP approximation were usually worse than the ones obtained in the FOM with the original modes from Hogan's study. In section 6.5 we discuss the implications of these results.

6.4 Chameleon Mode

In the previous FOM study there was no guarantee of improvement in the obtained trajectory by considering new modes in a previously used family. It needs to be clarified that, even if in the previous study there was no global score function implemented, this effect could be so big that it could be observed with the naked eye. An illustrative example would be the following. A family containing only the sticking mode $F_1 = \mathbf{S}$ would greatly outperform the extended family $F_2 = \mathbf{S} \cup (L, R, S, \overset{N-2}{\cdots}, S)$. On the other side a further extension $F_3 = F_2 \cup \mathbf{R}_1$ would outperform the two previous families. This behaviour can be shown in Figure 14.

The performance decrease experienced in F_2 can be explained by the fact that the FOM optimizer can choose optimal trajectories that will never be able to be applied. For example, the optimal control solution at a certain optimization iteration can be to choose the second mode of F_2: start with an L state for the first time step to then choose a R step and finally choose an S step for the rest of the finite time horizon prediction. The capability of sliding left and right should allow the pusher to take more aggressive turns by sliding to one of the slider edges and then slide back to the center to properly track the straight line. The main problem is that the control law for the controller is to execute the only th first step of the optimal trajectory, so the system will be in an L state for one time step and then the control law will be recomputed. As no mode starts with a sliding right state the slider can never go back to the central position.

The conclusion these experiences led us to was that, using Bersekas' terminology, some families did not have the sequential improvement property. What's more, the results seemed to indicate that composition or expansion of families that apparently had sequential improvement could lead to losing the property.

Solving the MPC using the FOM approach has similarities with the roll-out algorithm presented in their study: Each of the modes can be considered a heuristic and the strategy of recomputing the control policy

Figure 14: Local cost over time for families F_1, F_2, and F_3 obtained following a straight line on the x axis and using an initial condition $\mathbf{x}_0 = \left(0, 0.05, -\frac{\pi}{6}, 0\right)^T$.

at a certain frequency by using the same heuristics can be viewed as a roll-out-like technique. In fact, we could easily reformulate any family of our FOM approach to behave as their roll-out by increasing the finite horizon by one step and substitute each of the original modes by a triplet of copies preceded with an S, L, and R state respectively. We decided not to use an exact roll-out because, as we will show in this section, we were able to find a strategy to guarantee a sense of improvement without having to increase the number of MPC sub-problems to solve by a multiples of the number of possible states.

It is important to note that the concepts of sequential consistency or optimality cannot be applied to our case without some previous modifications, because of the existence of a receding finite time horizon, but for every FOM problem we can formulate an associated Full Horizon Family of Modes (FFOM) problem as follows:

Definition 6.1. FFOM:
Given a FOM problem defined by a finite horizon of N steps, a set H of possible hybrid states, a cost function J, and a family $F^{(0)} = \{M_1^{(0)}, ..., M_k^{(0)}\}$ containing k modes, where $M_i^{(0)}$ is a general mode described as $M_i^{(0)} = (s_{i,1}, s_{i,2}, ..., s_{i,N})$, with $s_{ij} \in H$. Furthermore, consider a cost function $J_k = \sum\limits_{i=k}^{N} \left(\bar{\mathbf{x}}_i^T \mathbf{Q} \bar{\mathbf{x}}_i + \bar{\mathbf{u}}_i^T \mathbf{R} \bar{\mathbf{u}}_i\right)$ (defined using the $\mathbf{Q}$ and $\mathbf{R}$ matrices from the FOM cost function J) and the prefixes of each of the modes of the family $M_i^{(j)} = (s_{i,1}, ..., s_{i,N-j})$. Then the associated FFOM problem is defined iteratively as follows:

1. Solve the FOM-based MPC problem defined by $F = F^{(0)}$ and $J = J_0$, using x_0 as the initial condition.

2. Use the controller $\bar{\mathbf{u}}_0$ from the optimal solution of the problem solved in the previous step to compute the value of $\bar{\mathbf{x}}_1$. Under an ideal world assumption (that is, assuming that our motion equations model

is accurate) the value of $\bar{\mathbf{x}}_1$ will correspond to the same value obtained in the optimal trajectory of the MPC.

3. Replace F by $F^{(1)} = \{M_1^{(1)}, ..., M_k^{(1)}\}$, J by J_1, and the initial condition by $\bar{\mathbf{x}}_1$ and go to step 1, solving this time an MPC problem of $N-1$ steps.

4. Repeat until the horizon has only one step.

This formulation represents how the FOM problem would behave be if we had enough computation power to cover a finite horizon that reaches the desired final time step. Using it, we can study the FFOM problem associated to the example families presented at the beginning of this section. For simplicity, we will consider the realized states and control inputs to describe the considered trajectories.

Example 6.2. F_1 is sequentially consistent in its FFOM version and assuming an ideal world:

Consider that at a certain control iteration there's a full horizon of N steps to reach the end of the trajectory. The only considered mode, $\mathbf{S}$, get the optimal path to reach the goal while keeping the sticking contact at all times, obtaining a sequence of position states $\mathbf{x}^{\mathbf{S}}(\mathbf{x}_0) = (\mathbf{x}_{1,1}, \mathbf{x}_{1,2}, ..., \mathbf{x}_{1,N})$ and control inputs $\mathbf{u}^{\mathbf{S}}(\mathbf{x}_0) = (\mathbf{u}_{0,1}, \mathbf{u}_{1,1}, ... \mathbf{u}_{1,N-1})$ for the optimal path from the initial position to the goal. Then, the controller will execute $\mathbf{u}_0$ for one control step and the pusher slider system will move to $\mathbf{x}_1$. The only heuristic it considers in the next optimization step is to stick until reaching the goal, so the optimal trajectory obtained will be $\mathbf{x}^{\mathbf{S}}(\mathbf{x}_1) = (\mathbf{x}_{1,2}, ..., \mathbf{x}_{1,N})$ and $\mathbf{u}^{\mathbf{S}}(\mathbf{x}_1) = (\mathbf{u}_{1,1}, ... \mathbf{u}_{1,N-1})$, being sequentially consistent.

Example 6.3. On the same conditions, F_2 is not sequentially consistent nor improving:

Let's call the second mode of family $\mathbf{LR}_1$. If this mode is chosen at a certain time step, giving an optimal trajectory $\mathbf{x}^{\mathbf{LR}_1}(\mathbf{x}_0)$ and control $\mathbf{u}^{\mathbf{LR}_1}(\mathbf{x}_0)$. In the next step, a mode $\mathbf{R}_1$ is required in the family to consider the continuation of the previous optimal trajectory, which consists in sliding right and then sticking until the end of the trajectory. As it is not the case, the new FFOM step will not consider the previously optimal trajectory and, if no relation can be proven between the trajectories considered in the current step and the ones considered in the previous step, neither consistency nor improvement cannot be guaranteed.

Consider that, at a given step of a FFOM problem, the initial position is given by $\mathbf{x}_0$ and the optimal trajectory is determined by the mode $M_1^{(0)} = (s_{1,1}, s_{1,2}, ..., s_{1,N})$. If the family contains a mode whose maximal prefix corresponds to the maximal suffix of $M_1^{(0)}$, $M_2^{(0)} = (s_{1,2}, ..., s_{1,N}, s_{2,N+1})$ then sequential improvement will be guaranteed at the next iteration step. That is, if the trajectory provided by M_1 at the first step is $\mathbf{x}^{M_1^{(0)}}(\mathbf{x}_0) = (\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_N)$ and $\mathbf{u}^{M_1^{(0)}}(\mathbf{x}_0) = (\mathbf{u}_0, \mathbf{u}_1, ..., \mathbf{u}_{N-1})$ and the one provided by M_2 at the second step is $\mathbf{x}^{M_2^{(1)}}(\mathbf{x}_1) = (\mathbf{x'}_2, ..., \mathbf{x'}_N)$ and $\mathbf{u}^{M_2^{(1)}}(\mathbf{x}_1) = (\mathbf{u'}_1, ..., \mathbf{u'}_{N-1})$, then

Lemma 6.4. $J_1(\mathbf{x}^{M_2^{(1)}}(\mathbf{x}_1), \mathbf{u}^{M_2^{(1)}}(\mathbf{x}_1)) = J_1(\mathbf{x}^{M_1^{(1)}}(\mathbf{x}_1), \mathbf{u}^{M_1^{(1)}}(\mathbf{x}_1))$:

As $M_2^{(1)}$ is the maximal suffix of $M_1^{(0)}$ the continuation of the trajectory provided by M_1 on the first step is feasible in the mode M_2 during the second step so, as the trajectory given by M_2 is optimal, $J_1(\mathbf{x}^{M_2^{(1)}}(\mathbf{x}_1), \mathbf{u}^{M_2^{(1)}}(\mathbf{x}_1)) \leq J_1(\mathbf{x}^{M_1^{(1)}}(\mathbf{x}_1), \mathbf{u}^{M_1^{(1)}}(\mathbf{x}_1))$

If we assume $J_1(\mathbf{x}^{M_2^{(1)}}(\mathbf{x}_1), \mathbf{u}^{M_2^{(1)}}(\mathbf{x}_1)) < J_1(\mathbf{x}^{M_1^{(1)}}(\mathbf{x}_1), \mathbf{u}^{M_1^{(1)}}(\mathbf{x}_1))$ *we can easily reach a contradiction. Let's define* $\mathbf{x}^* = (\mathbf{x}_1, \mathbf{x'}_2, ..., \mathbf{x'}_N)$ *and* $\mathbf{u}^* = (\mathbf{u}_0, \mathbf{u'}_1, ..., \mathbf{u'}_{N-1})$. *This is a feasible trajectory in* $M_1^{(0)}$ *but* $J_0(\mathbf{x}^*, \mathbf{u}^*)$ *can be expressed as:*

$$J_0(\mathbf{x}^*, \mathbf{u}^*) = \mathbf{x}_0^T \mathbf{Q} \mathbf{x}_0 + \mathbf{u}_0^T \mathbf{R} \mathbf{u}_0 + J_1(\mathbf{x}^{M_2^{(1)}}(\mathbf{x}_1), \mathbf{u}^{M_2^{(1)}}(\mathbf{x}_1))$$

and we can conclude that $J_0(\mathbf{x}^*, \mathbf{u}^*) < \mathbf{x}_0^T \mathbf{Q} \mathbf{x}_0 + \mathbf{u}_0^T \mathbf{R} \mathbf{u}_0 + J_1(\mathbf{x}^{M_1^{(1)}}(\mathbf{x}_1), \mathbf{u}^{M_1^{(1)}}(\mathbf{x}_1)) = J_0(\mathbf{x}^{M_1^{(0)}}(\mathbf{x}_0), \mathbf{u}^{M_1^{(0)}}(\mathbf{x}_0))$, reaching a contradiction.

Let's now consider an extension state E. That is, a default state to apply mode extensions. We can then define a mode suffix operation as follows.

Definition 6.5. Mode Suffix Operation:

Given a default extension state E, applying the suffix operation to mode $M_1 = (s_1, s_2, ..., s_N)$ results in $M_2 = suff(M_1) = (s_2, ..., s_N, E)$

This definition allows us to announce the following corollary from lemma 6.4.

Corollary 6.6. *A family F closed by a suffix operation is sequentially improving in the FFOM problem:*

At any step i of the approximated MPC, a certain mode M will provide the optimal trajectory. As $suff(M) \in F$ we can use lemma 6.4 to guarantee that the trajectory provided by $suff(M)$ in step $i + 1$ is as good as the one currently provided by M. As the approximated MPC considers the optimal trajectory between all the considered modes, the computed trajectory at $i + 1$ will be at least as good as the one computed in i.

Example 6.7. On the same conditions, F_3 is sequentially improving:

F_3 is closed by the suffix operation defined using S as the extension state, so we can use corollary 6.6 to guarantee sequential improvement in the FFOM.

We considered important to guarantee sequential improvement for the FFOM problem even taking into account that our problem is not in the ideal world (modelling errors and different frequencies for the optimization problem and the controller can lead the state obtained when applying u_0 to $\bar{x}_0$ to be different from x_1) and that we don't have a full horizon. As we showed in the examples above, the simulation results were in agreement with these theoretic results as the families that implicitly guaranteed sequential improvement in their FFOM version tended to outperform the ones that did not.

In order to guarantee it we decided first decided to only use families that were closed by the suffix operation obtained using **S** as the extension mode. The problem with this approach is that for any mode $M = (s_1, ..., s_N)$ you add to a family F you need to include all of its possible **S**-extended suffixes. Different modes required adding a different number of suffix modes. For example, an **S** mode required no extra modes (as it is itself closed by the suffix operation) and for a $\mathbf{R}_n$ one you only need to extend the family with an **S** mode and all the $\mathbf{R}_i$ modes for $i < n$. Even taking this into account, we considered adopting this strategy constrained the problem of family design more than we wanted to and added more extra nodes that was strictly necessary, so we decided to adopt another strategy. This translated into changing from static families to dynamical ones.

The strategy we used consisted in including a dynamic mode that we called chameleon mode. If, at a certain step i, the optimal FOM mode is mode M, the chameleon mode is defined in step $i + 1$ as $suff(M)$ (using, as before, **S** as the extension state). This guaranteed sequential improvement by applying lemma 6.4 in an analogue way as in corollary 6.6.

The chameleon mode functionality was added in the simulation interface and its performance was tested in simulation. As we will show in section 6.5, implementing this functionality resulted in an improvement of the sequential behaviour of the FOM algorithm. Being able to safely add arbitrary new modes to the FOM family without fear of losing its sequential properties allowed us to develop an strategy to learn a good family set-up for a general trajectory tracking problem, as will be shown in chapter 7.

Given the good properties of the chameleon mode, we decided to include it in the FOM formulation so, from now on,we will use FOM to refer to the family of modes with chameleon mode implementation, unless specified otherwise.

6.5 Simulation Results

In order to compare FOM and MIQP we considered necessary to use the same finite horizon length in both formulations to have a fair comparison. As explained in section 6.3, the computational cost was much higher than we expected. The maximum finite horizon length we could afford to run simulations in a reasonable time was 5. This value of the finite horizon was too short to control most of the non-trivial trajectories properly. Figure 15 shows the trajectory obtained from MIQP while trying to follow a circular trajectory and it can be seen that it clearly diverges.

Figure 15: Trajectory obtained when tracking circular trajectory defined by pushing perpendicularly to the slider with $p_y = 0.0249$. Initial position $\mathbf{x}_0 = (0.0, -0.05, -0.1, p_y)^T$.


Figure 16: Trajectory obtained while tracking the same nominal trajectory as Figure 15 starting at the same initial condition and using a clustered MIQP solver consisting in 5 clusters of 7 steps.

We tried to track the same trajectories using our implementation of MIQP with 5 clusters of 7 steps. This allowed us to compare it with a FOM problem of 35 steps. Using the original family from Hogan's work (formed bye the modes $\mathbf{S}$, $\mathbf{R}_1$, and $\mathbf{L}_1$) we can observe that, for non-trivial trajectories, the FOM approach outperforms the clustered MIQP one. As can be seen in Figure 18 the clustered MIQP does not even converge to the nominal trajectory.

The results obtained with MIQP and clustered MIQP imply that we will not be able to compare the FOM solvers with an ideal solution. As will be seen in future sections, the original modes from Hogan's study are specially versatile and can track a wide range of trajectories. Because of that, in chapter 7 we will use them as a comparison model to strive for.

Figure 17: Trajectory obtained while tracking the same nominal trajectory as Figure 15 starting at the same initial condition and using a FOM solver with a finite horizon of 35 steps and the original modes from Hogan's study.
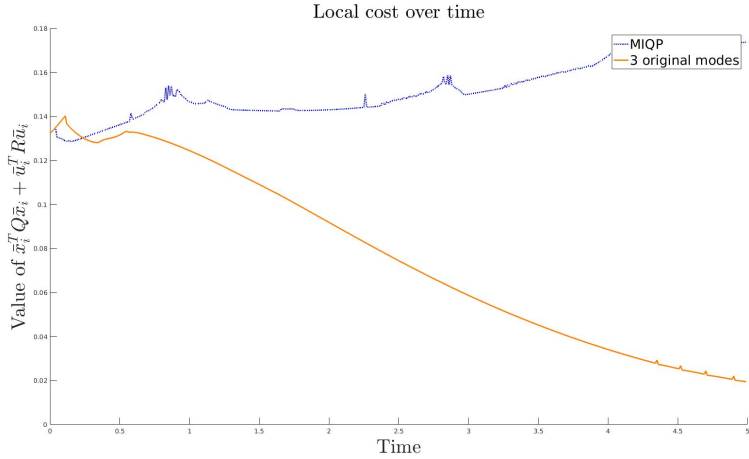

Figure 18: Local cost over time for the trajectories shown in Figures 16 and 17.

Finally, we tested our chameleon mode implementation by comparing the performance of families of modes that are not sequentially consistent in their full horizon formulation with and without the chameleon mode with satisfactory results. In Figure 19 we show the example presented in section 6.4. All the simulation parameters are kept the same except for the usage of the chameleon mode. As can be observed, the results correspond with our expectations: for the family F_2 the chameleon mode implementation successfully tracks the trajectory while the implementation without it failed to do it, as we showed previously.

Figure 19: Local cost over time for families F_1, F_2, and F_3 obtained following a straight line on the x axis and using an initial condition $\mathbf{x}_0 = (0, 0.05, -\frac{\pi}{6}, 0)^T$.

7. Mode Selection Protocol

In the previous sections we explained how we addressed some of the limitations or mistakes in the FOM formulation. We also tried to provide an ideal solution to strive for and we concluded that it was computationally intractable. At this point of our study, focused on adressing our next objective: to provide a selection protocol for the modes used in the FOM formulation.

Our main goal was to obtain a simple and general procedure to select modes that would track a wide variety of trajectories. We were specially interested in keeping the protocol simple because our ultimate goal was to expand the FOM usage into the research community, so we needed an out-of-the-box general method that would not require much extra set-up steps depending on the particular manipulation system studied. In order to formalize the mode selection, we decided to represent it as an optimization problem.

We considered multiple optimization criteria and relations to study and we had to adopt a set of assumptions to be able to obtain some initial results. We first tried to learn the modes to follow a single fixed trajectory. Each optimization problem was parametrized by a fixed finite time horizon (characterized by a number of steps N) and each family had a fixed number of modes K. The initial condition $\bar{x}_0$ was assumed to be a random variable following a certain probability distribution. The cost of the control trajectory C was thus also a random variable of unknown distribution parametrized by the optimization problem parameters. Specifically, given a set of K modes $\{m_1, ..., m_k\}$, $C \sim D(m_1, ..., m_k)$ for a given distribution law D. Ideally, we would like to choose the the set of modes that minimizes $\bar{C}$. We considered that this stochastic approach would provide us versatile modes that would work under a wide variety of initial conditions, rather than over-fitting a specific set of them.

We made three main assumptions:

Assumption 7.1. Local Cost Sufficiency: *Optimizing the modes to reduce the local costs obtained while solving each MPC iteration would be sufficient to obtain a low global cost.*

Assumption 7.2. Locally Invariant Trajectories: *Given a local inertial frame F centered on the slider center of mass moving and rotating with it and given an infinite objective trajectory $C(t)$ for $t \in (-\infty, \infty)$, then the $C(t)$ would be time invariant in frame F (for example an infinite circle loop or an infinite straight line).*

Assumption 7.3. Single Step Sufficiency: *A single control iteration of the trajectory tracking problem was sufficient to learn which modes where optimal, if the appropriate distribution for $\bar{x}_0$ was chosen.*

The motivation and justification for these assumptions will be explained in 7.1, as well as some other implicit assumptions contained in it. In 9 we will discuss some of the studies we would have conducted and how we would have modified these assumptions if we disposed of more time.

7.1 Structure Study And Systematic Approach

Given a finite horizon of N steps and H states and considering M the set of possible modes obtained by concatenating N of these hybrid states, then $|M| = H^N$. Let the set of possible families containing K of these modes be F. Then $|F| = \binom{|M|}{K} = \frac{H^N!}{K!(H^N-K)!}$. Even with a low number of states this number grows quickly as N and K increase.

A brute force approach to the optimization problem would prove to be computationally intractable and, without any further assumption or knowledge of the structure of the problem, we would have to take this brute force approach. There were some correlation effects between modes that we wanted to capture. For example, a mode L_1 usually needs some other mode in the family that allows to slide right, as R_1. Otherwise if the slider slides to the left at a certain time step, the pusher will never be able to recover its initial position. In section 6.5 we will show that, as expected, once a pusher-slider system reaches a configuration with the pusher near one of the slider edges it is very difficult that the controller moves the pusher the other side of the slider as a result of the FOM optimization. This configuration allows to track steep turns in one direction, but will probably fail to control at all trajectories that turn in the other direction. As we will explain below, the combination of the assumptions mentioned in the previous section provided a way to capture this correlations considering the modes individually.

The local cost sufficiency assumption was listed as a new assumption, but it was already implicitly taken when deciding to use a finite horizon MPC formulation to control the pusher-slider system. The mode selection will determine the performance of our local FOM solver and the relation between local MPC performance and overall performance and the validity of this assumption was already discussed in 6.1.

The locally invariant trajectories assumption motivation will be used in combination with section 7.3, where we will use them to approximate general trajectories. It should be noted that for a trajectory to be locally invariant, it has to be defined by the same hybrid state at every instant. Furthermore, the only state which allows infinite locally invariant trajectories is the sticking state (in the other two the pusher would keep slipping until losing contact with the slider). So, the only trajectories we will consider are the ones which keep the sticking state at every instant. In the next section we will characterize the considered trajectories.

Remark 7.4. Note that in all the considered trajectories, the nominal mode is **S**. We will use this fact at the end of this section.

The main motivation for the single step sufficiency assumption has to be understood in combination with the previous assumption. The locally invariant property implies that the only difference between an MPC step and another is their initial condition random variable $\bar{\mathbf{x}}_i$. Consider a distribution law **X** that focusses more on covering a wide range of possible values of any $\bar{\mathbf{x}}_i$ rather than on their frequency as, for example an appropriate uniform distribution. We can then optimize the modes to reduce the expected value of $C(\bar{\mathbf{x}}_0)$ of a single controlling step tracking problem, where $\bar{\mathbf{x}}_0 \sim \mathbf{X}$. Optimizing the mode selection this will simultaneously reduce the local costs expected value for all $\bar{\mathbf{x}}_i$ implicitly assuming that each of these values are equally distributed.

This implicit assumption is clearly false. The fact that some trajectories converge is already a counterexample of it. We consider assuming a uniform to describe all $\bar{\mathbf{x}}_i$ is a conservative approach, as will consider rare big disturbances (that can usually only happen with an external perturbation or a bad initial position) to be equally likely to appear at any time step as small disturbances characteristic of proper control around the objective trajectory. We expected this approach would be sufficient to provide a mode selection robust enough to track a wide range of trajectories and initial conditions.

These assumptions result in a great cost reduction and simplification of the learning algorithm because the optimization problem can be now be easily approximated as a function of samples of $\bar{\mathbf{x}}_0$:

Corollary 7.5. The optimization problem to select the FOM modes can be expressed as:

$$\min_{m_1,\dots,m_K\in M} \mathbb{E}\big(C(\bar{\mathbf{x}}_0; m_1, \dots, m_K)\big) = \min_{m_1,\dots,m_K\in M} \mathbb{E}(\min_{i=0}^{K} C_{m_i}(\bar{\mathbf{x}}_0)) \approx \min_{m_1,\dots,m_K\in M} \frac{\sum_{j=0}^{E}(\min_{i=0}^{K} C_{m_i}(\bar{\mathbf{x}}_j))}{E}$$

Where E is the number of samples of $\bar{\mathbf{x}}_0$ we consider.

Where C_i is the cost associated to mode m_i in the MPC optimization. This further implies that solving the MPC problem for all the possible sets of K modes is no longer necessary, as we announced at the beginning of this section because, if $C_i(\bar{\mathbf{x}}_0)$ is computed for every mode $m_i \in M$, then $C(\bar{\mathbf{x}}_0; m_1, \dots, m_K)$ is defined for any set of K modes by corollary 7.5.

We were concerned that the initial condition would probably have a lot of effect in the overall cost of the mode-condition pair, because higher or smaller disturbances would lead to higher or smaller overall cost (independently of the mode chosen) and the simulation results from section 6.5 confirmed it, as can be seen in Figure 20. We wanted to reduce the impact of the initial condition in the range of possible values of the local costs to obtain comparable experiments, so we decided to standardize the results obtained within each experiment. Consider a specific sample $\bar{\mathbf{x}}_j$ of the random variable $\bar{\mathbf{x}}_0$, then we can compute the statistical mean of the local costs over all the modes $\bar{C}^{(j)} = \frac{\sum_{i=0}^{K} C_i(\bar{\mathbf{x}}_j)}{K}$ and $\sigma^{(j)} = \sqrt{\sum_{i=0}^{K} \frac{(C_i(\bar{\mathbf{x}}_j) - \bar{C}^{(j)})^2}{K-1}}$ and we can use them to compute the normalize costs $\hat{C}_i(\bar{\mathbf{x}}_j) = \frac{C_i(\bar{\mathbf{x}}_j) - \bar{C}^{(j)}}{\sigma^{(j)}}$. Figure 21 shows the distribution of the local costs over modes after the normalization. As this values were more comparable, we minimized the normalized costs instead of the original ones.

The algorithmic cost reduction obtained from corollary 7.5 was not enough to face the full optimization problem. We no longer had to study the full space of families F, but we still had no alternative to explore
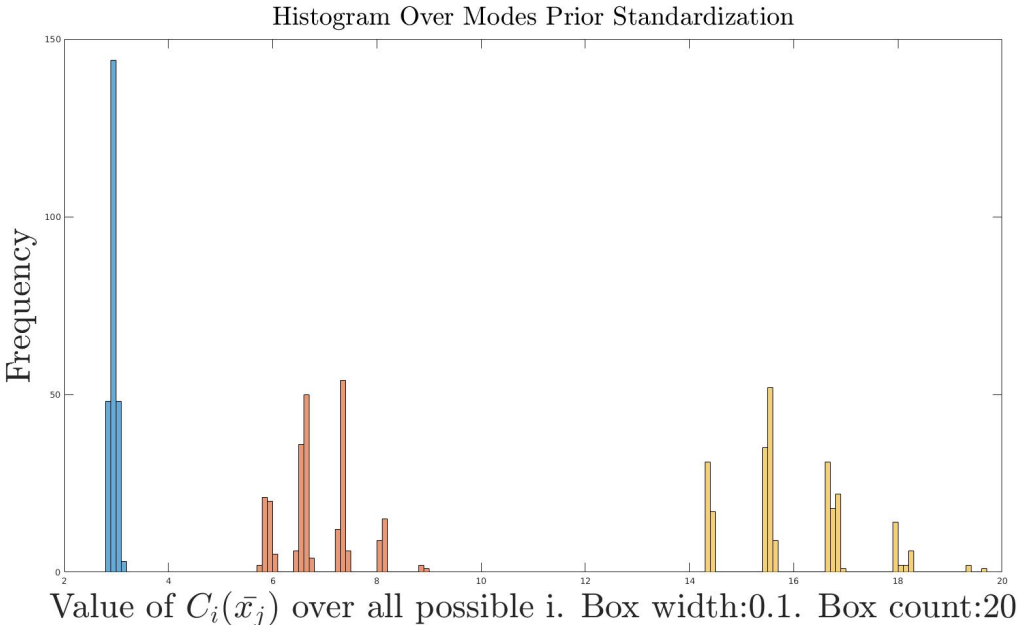
Figure 20: Example histogram of $C_i(\bar{x}_j)$ of the local costs obtained while tracking a circle for 1000 initial conditions and 253 possible modes. The histogram is taken by fixing an initial condition $\bar{x}_j$ and evaluating the values of $C_i(\bar{x}_j)$ for all possible i. Each color represents a different initial condition. The costs are clearly non-comparable
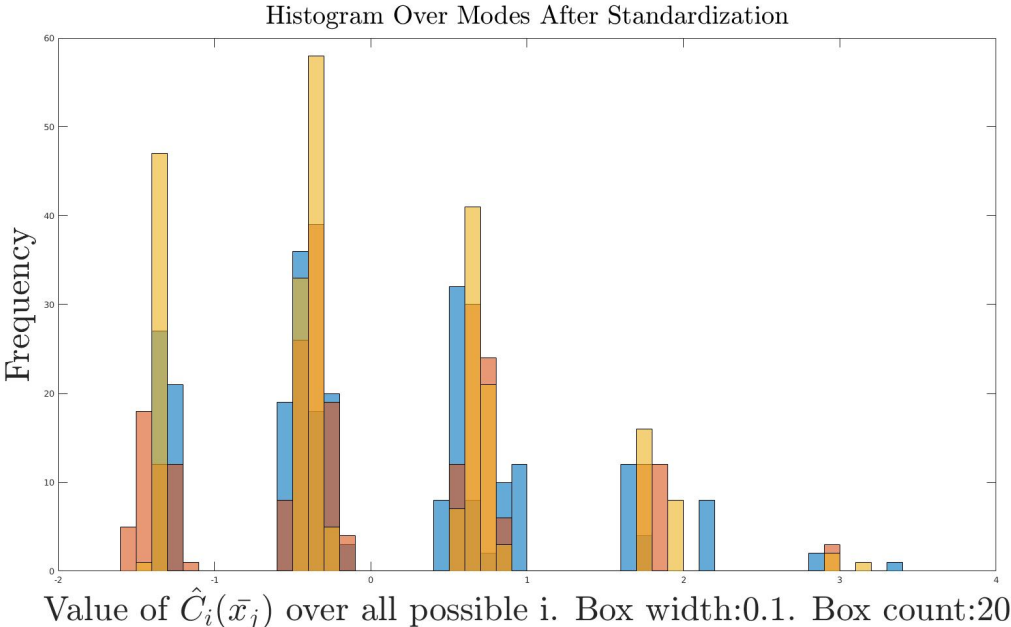

Figure 21: Example histogram of $\hat{C}_i(\bar{x}_j)$ for the same simulation set-up as in Figure 20. The costs are now comparable

the full space of modes M, which grows exponentially with the length of the finite horizon. Inspiring ourselves in the original modes from Hogan's study, that concentrate all their variability on the first step, we devised the following protocol:

1. Consider short horizons that allow to compute $\hat{C}_i$ for all the possible modes fast.

2. Gather a number samples of $\hat{C}_i$ relatively big compared to the number of possible modes.

3. Select the modes that minimize the cost function in the small horizon environment.

4. Extend the families to a longer horizon with simple S state tails.

As we will show in section 7.3, this approach is sufficient in our pusher-slider system to track complex trajectories, but in section 9 we discuss an iterative procedure to extend the considered finite horizon.

Choosing the set of K modes that minimizes the desired expression can depend exponentially on K, so we used a greedy upper bound approximation to compute it. The greedy approach was defined iteratively as follows:

1. Consider $m_1 = \mathbf{S}$ (it is the nominal mode, so we always need to include it in the family) and $C_1(\bar{\mathbf{x}}_j) = C_{m_1}(\bar{\mathbf{x}}_j)$

2. Select m_{i+1} to be the one that minimizes $\dfrac{\sum\limits_{j=0}^{E} \min(C_i(\bar{\mathbf{x}}_j), C_{m_{i+1}}(\bar{\mathbf{x}}_j))}{E}$

3. Let $C_{i+1}(\bar{\mathbf{x}}_j) = \min(C_i(\bar{\mathbf{x}}_j), C_{m_{i+1}}(\bar{\mathbf{x}}_j))$

This algorithm provides an upper bound for $\min\limits_{m_1,...,m_K \in M} \dfrac{\sum\limits_{j=0}^{E} (\min\limits_{i=0}^{K} C_{m_i}(\bar{\mathbf{x}}_j))}{E}$ we implemented the algorithm and we determined the lower bound of $\bar{C}(\bar{\mathbf{x}}_0; m_1,...,m_K)$ by computing the statistical mean over the minimum of all the modes: $\dfrac{\sum\limits_{j=0}^{E} (\min\limits_{i=0}^{|M|} C_{m_i}(\bar{\mathbf{x}}_j))}{E}$. Figure 22 shows the evolution of the cost provided by our greedy algorithm as more modes are added to the family. It can be seen that the upper bound quickly converges to the lower one.

In the following section we will apply this procedure to linear and circular trajectories.

7.2 Simulation Results. Simple trajectory learning

For our initial study, we tried to learn the best modes for single locally invariant trajectories. As proposed in the previous section, we started by obtaining 1000 samples of $\bar{\mathbf{x}}$ from a uniform distribution in a cube in $\mathbb{R}^4$. The $\bar{x}$ and $\bar{y}$ disturbances on the x and y position of the slider center of mass were bounded between 0.5 and 0.5 meters. The slider angle disturbance $\bar{\theta}$ was bound between $-\dfrac{\pi}{4}$ and $\dfrac{\pi}{4}$ and the value of p_y could lay in any point on the slider surface. For each initial condition, the tracking problem was solved for all the possible modes with a small finite horizon of 5 steps. The number of possible modes was then 243 and each pair mode-sample had an associated single step FOM tracking problem. The cost functions $C_i(\bar{x}_j)$

Plot Of Normalized Total Cost Over Number Of Modes

Figure 22:   Plot of our greedy approximation to the minimum value of $\bar{C}(\bar{\mathbf{x}}; m_1, ... m_K)$ obtained by using families with K modes. The horizontal bar shows the optimal lower bound obtained when all the families are used. Its exact value is -1.5819.
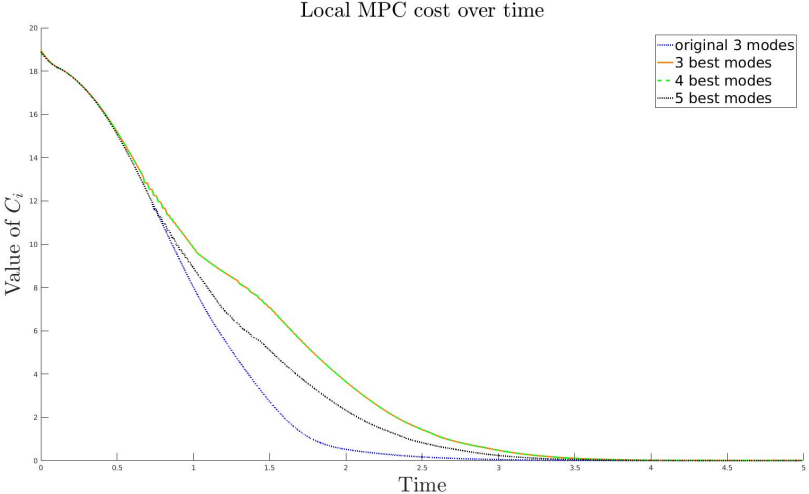
Local MPC cost over time

Figure 23:   Comparison between the original FOM from Hogan's study and the families obtained by applying the learning algorithm to a the linear trajectory $\mathbf{x}^* = (vt, 0, 0, 0)^T$ $\mathbf{u}^*(t) = (v, 0)^T$. The initial condition is $\bar{\mathbf{x}}_0 = (0, .05, \frac{\pi}{6}, 0)^T$. The vertical axis correspond to the optimal value of the finite horizon MPC over time.

were computed for each mode-initial condition pair and were normalized following the procedure described in the previous section.

Figure 24: Trajectory obtained when tracking the system described in Figure 23 with a family containing the best 4 line modes learned using this section's procedure.

We then used the greedy algorithm we developed to compute and store our approximate ordering of modes and we extend them to a longer horizon of 35 steps and we used the k best modes to track base trajectories.

We applied this procedure to several circular trajectories and a horizontal line trajectory defined by $\mathbf{x}^* = (vt, 0, 0, 0)^T$ $\mathbf{u}^*(t) = (v, 0)^T$ for a certain pusher velocity v. The results for linear trajectories are shown in Figures 23 and 24. All the learned families properly converge approximately at the same time and have a similar behaviour.

The circular trajectories were imposed by continuously pushing perpendicular to the slider side ($\mathbf{u}^*(t) = (v, 0)^T$) at a distance d of the center of the slider side ($p_y(t) = d$). Applying the motion equations derived in section 5.5 and using the following definitions allows us to nicely compute the analytical nominal trajectory: Let $A = \dfrac{v}{c^2 + p_x^2 + d^2}$, $B = (c^2 + p_x^2)$, and $C = p_x d$. Then $\dot{\theta}(t) = -A$ so, assuming $\theta(0) = 0$, then $\theta(t) = -At$. Then:

$$\begin{pmatrix} \dot{x}(t) \\ \dot{y}(t) \end{pmatrix} = \frac{A}{d} \begin{pmatrix} cos(\theta(t))B - sin(\theta(t))C \\ sin(\theta(t))B + cos(\theta(t))C \end{pmatrix} \tag{27}$$

Which can be integrated to obtain:

$$\begin{pmatrix} x(t) \\ y(t) \end{pmatrix} = \begin{pmatrix} x(0) \\ y(0) \end{pmatrix} + \frac{A}{d\dot{\theta}} \int_0^T \dot{\theta} \begin{pmatrix} cos(\theta(t))B - sin(\theta(t))C \\ sin(\theta(t))B + cos(\theta(t))C \end{pmatrix} dt$$

considering $x(0) = y(0) = 0$, we obtain:

$$\begin{pmatrix} x(t) \\ y(t) \end{pmatrix} = \frac{-1}{d} \begin{pmatrix} sin(\theta(t))B + cos(\theta(t))C - C \\ -cos(\theta(t))B + sin(\theta(t))C + B \end{pmatrix} \tag{28}$$

Also, from this expression we can easily compute the curvature associated with this curve, that we will use in the next section:

$$k = -sign(d)\frac{1}{R} = -sign(d)\frac{d}{\sqrt{B^2 + C^2}} \tag{29}$$

Remark 7.6. We will use the common mathematical notation for the curvature sign. That is, a positive curvature means that the unit tangent vector rotates counter-clockwise as a function of the parameter along the curve.

Figures 25 and 26 show the results of tracking one of the learned circular trajectories. As in the case of the linear trajectory, all the modes have similar behaviour. This is consistent with the properties shown in Figure 22. The normalized cost the modes are minimizing quickly approaches the optimal in few iterations of the greedy algorithm and afterwards there's minimal improvement until reaching the optimal value.
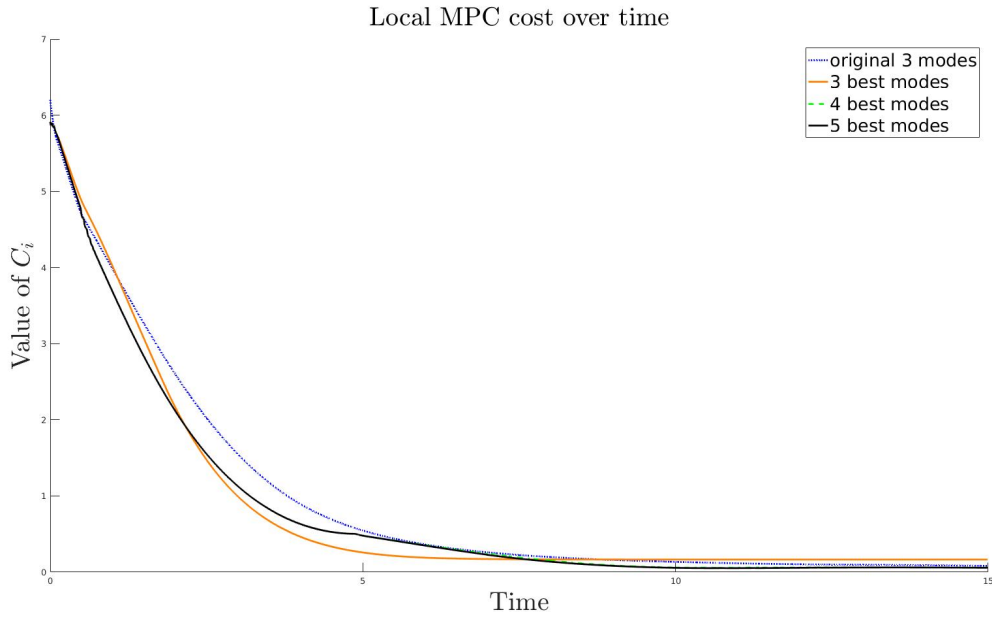

Figure 25: Analogue to Figure 23 for a circle trajectory with a internal disturbance. The circular trajectory corresponds to the one obtained by setting $d = 0.0249$. The initial condition is $\bar{x}_0 = (0.0, -0.05, -0.1, 0.0249)^T$

Figure 26: Trajectory obtained while tracking the system described in Figure 25 with the family formed by the 4 best modes learned for this circle.

7.3 Dynamic FOM. Theory and results

The assumption of local invariance for the objective trajectories limited the range of application of the FOM approach as it was presented until now. In section 7.1 we mentioned that we had a motivation to use locally invariant trajectories before obtaining the simulation results from the previous section. This motivation was to train the FOM method on several base locally invariant trajectories that could describe most of the situations the tracking problem would face (For the simple contact pusher-slider system this consisted in a straight line and several circular trajectories of different curvatures). Then we wanted to learn the best k modes for every base trajectory and then approximate complex trajectories them as a piece-wise compositions of base ones. At each controller step we would use the best modes for each base trajectory depending on the base trajectory approximation of the current step.

For simplicity, we decided to approximate the motion of the finite horizon by the local geometry of the instantaneous position in the nominal trajectory. We expected that a sufficient control frequency would allow the system to properly track the general trajectory even if its local properties undertook steep temporal changes. For the pusher-slider system mentioned before, the base trajectory was chosen by numerically computing the curvature of the objective trajectory at every control step and using the base trajectory with the most similar curvature value.

We would then introduce dynamic families to the FOM problem. At each control step, the family used to track the trajectory would consist on the k best modes for the base trajectory that approximates the local geometry of the objective trajectory (together with the chameleon mode) to solve the MPC problem on that instant. The chameleon mode was shared between the dynamical families to ensure that the sequential

Figure 27: DFOM with 20 curvatures using the 4 best modes for each base trajectory. Finite horizon of 35 steps.


Figure 28: DFOM with 20 curvatures using the 4 best modes for each base trajectory. Finite horizon of 15 steps.

consistency property was not lost.

The main motivation for this method was to reduce the online cost of the FOM approach at the expense of more offline computations. We expected that choosing the best family dynamically would reduce the number of modes per family required to obtain good results in a broad range of trajectories. Having n modes in a family implies solving n convex optimization problems at every control step, so we wanted to keep it as low as possible.

The learning algorithm was executed on 10 pairs of curvatures (for each curvature k, we considered the complementary curvature $-k$) generated by setting the d parameter uniformly from $-\dfrac{3b}{4}$ to $\dfrac{3b}{4}$, where b is the length of the side of the slider. Using equation 29, the associated curvatures were obtained.

Generating general feasible trajectories for the controller to stabilize is a complex problem in itself. We could use MPC to compute an optimal path to a certain destination using the distance to the destination as the cost function but this would have two important drawbacks. The first is that it would probably follow straight lines whenever possible, providing simple trajectories to stabilize. The second is that with MPC we can only obtain a discrete trajectory. We could approximate the obtained trajectory using splines, but we would not be able to guarantee feasibility of the trajectory defined by the state variable and the controller in the interpolated regions between control points.

The planning problem escaped the scope of this study and, in order to both simplify it and to check the robustness of the controller, we decided to create trajectories by simple concatenation of linear and circular trajectories. This simplified the planning problem and created punctual discontinuities in the position of the pusher (for the straight lines the pusher was constantly at the center of the slider side, and for the circular trajectories the pusher was constantly at the distance p_y associated to the circle and there were discontinuities at the changes between trajectories). We considered important to study whether the controller was capable to react to these temporally localized points of uncertainty were there was no feasible trajectory to immediately change from one trajectory to the next.

Remark 7.7. The circular trajectories used in the simulations of the DFOM had curvatures different than the ones used for the learning protocol, so the DFOM algorithm had to determine the base circle which was more similar to the nominal trajectory one.

With this formulation we were able to successfully track the general trajectory and we were able to reduce the required finite horizon to track properly the trajectory to 15 steps, as can be seen in Figure 28

8. Discussion and Conclusion

In the previous sections we obtained several positive results regarding the FOM algorithm. In this section we will expose them following the order in which they were presented in the thesis:

1. The results obtained in sections 6.2 and 6.4 are fundamental to approach FOM to the research community. The change in formulation that ultimately lead to the presentation of the nominal mode is fundamental to obtain the good results showed along chapter 7. We want to emphasize that, without it, all the controllers we tested on circular trajectories diverged as the one showed in Figure 12.

2. We would have liked to be able to compare the trajectories obtained by MIQP or by the clustered MIQP implementation we developed. In any case, showing the huge computational cost of MIQP in

our velocity based formulation is useful to get an idea of the urge of the manipulation community to find fast solutions to it, even if they are approximated, as FOM. While implementing it, we realized that its drawbacks are not only the lack of scalability as the number of variables increases. Even if the same number of variables is uses, if the geometry of the optimization problem changes the computational cost of the problem can increase several orders of magnitude, as we explained in sections 6.3 and 6.5.

3. The sequential study of the FOM families that ultimately led to the chameleon mode has proven to be critical towards increasing the understanding of the FOM mode and towards keeping it simple. If we had not studied and addressed the sequential weaknesses of FOM, we could not have ultimately devised the mode selection procedure, as adding modes to families would generally imply loss of sequential properties and, thus, worse performance. If we had decided to constraint our families to be sequentially improving (by, for example, enforcing closure with respect to a suffix operation) the size of the FOM families would probably have increased a lot in order to keep the sequential improvement.

4. While implementing the mode selection protocol we were able to better understand why the fast and exact MPC solving is still an open problem. In order to be able to compute the best modes for a certain trajectory in a reasonable time we had to take several layers of assumptions. Fortunately, we were finally able to automatically provide minimal families of around 4 modes that successfully tracked the trajectories they were designed for.

5. Showing that the DFOM is capable to track trajectories with punctual discontinuity is a great result toward proving the robustness and generality of the FOM approach, but we should be able to track other trajectories than compositions of circles and lines if we want to convince the manipulation community that the FOM can be a general controller.

If we look at these results as related to the objectives we presented in chapter 4, we can conclude that:

1. The properties and formulation of the FOM have been clearly improved thanks to the concepts of chameleon and nominal modes and all their related theoretical background.

2. We have not been able to compare the results of FOM with those of MIQP. We successfully compared FOM with clustered MIQP, but the fact that the second was worse doesn't allow us to have any ideal control trajectory to strive for.

3. We successfully provided a simple set-up for the FOM method that is capable to track base trajectories described by straight lines, circles or a combination of both.

4. Being able to track all these new trajectories also accounts to partially show the generality and robustness of the method, but we consider

In conclusion, we consider that all the results we obtained will bring the FOM method closer to the robotics manipulation community. However, we will continue this study to improve the FOM method further before presenting it to ISRR. In the next chapter we will expose the main lines of research we are currently studying and the list of studies/experiments we want to do before the symposium.

Figure 29: Graphical representation of a pusher-slider system following a general trajectory. The osculating circle has been computed on the current trajectory state point.

9. Current and Future Lines of Research

We want to further study the generality and robustness of the method, as we feel that our current results would not completely convince the manipulation community.

We are currently trying to expand the general trajectory tracking of the DFOM method by using a similar approach as in Hogan's work [14] [15]. Given a nominal trajectory $y^* = f(x^*)$ with $x^*(t) = t$ for simplicity, then the local curvature of the trajectory can be computed by:

$$k = \frac{f''(x)}{(1 + f'(x)^2)^{3/2}}, \tag{30}$$

Given the parametrization of the trajectory, we can assume $x(t)$ and $y(t)$ to be the ones defined by the trajectory, $\theta(t)$ to be $\theta(t) = f'(x)$ and we then assume that the trajectory is defined instantaneously by the osculating circle, so the value of d can be computed numerically so that equation 29 holds. As in the trajectory from section 7.3, the resulting equation has punctual discontinuities, so we cannot guarantee that the DFOM will be robust enough to track it until we finish the implementation and run the simulations.

At the same time, we extended the formulation of the single point pusher to define the motion equations of a line pusher. We are currently trying to reproduce the learning protocol used for the single point pusher with this new system. We will try to analyse the robustness of the FOM approach by using a single point pusher FOM controller to control a line pusher and vice versa.

On the learning protocol side, we will soon try an iterative method to learn the best modes for long horizons. The method starts as the procedure we presented in section 7.1. That is by learning the best k modes for a horizon of 5 steps. We then extend these k modes to a horizon of 10 steps by adding all the possible 3^5 tails to each of the best modes of the first iteration. We would then apply the same learning procedure to the new $253k$ modes, selecting again the k best.

This procedure could be repeated iteratively until reaching a long horizon of around 35 steps.

References

[1] Gurobi - mixed-integer programming (mip) - a primer on the basics. `https://www.gurobi.com/resources/getting-started/mip-basics`. Accessed: 2017-04-24.

[2] Gurobi optimizer description. `http://www.gurobi.com/products/gurobi-optimizer`. Accessed: 2017-04-24.

[3] G. Amontons. De la résistance causée dans les machines. *Mémoires de lAcadémie royale des Sciences*, pages 206–227, 1699.

[4] M. Anitescu and F. Potra. Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems. *Nonlinear Dynamics*, 14(93):231–247, 1996.

[5] Dimitri P. Bertsekas, John N. Tsitsiklis, and Cynara. Wu. Rollout algorithms for combinatorial optimization. *Journal of Heuristics*, 3(3):245–262, 1997.

[6] N Chavan-Dafle and A Rodriguez. Prehensile pushing: In-hand manipulation with push-primitives. *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015.

[7] R.W. Cottle, J.S. Pang, and R.E. Stone. The linear complementarity problem. *Academic Press*, 1992.

[8] C. A. Coulomb. Théorie des machines simples en ayant égard au frottement de leurs parties et à la roideur des cordages. *Mémoires de mathématique et de physique présentés à lAcadémie royale des Sciences*, 1781.

[9] Nima Fazeli, Roman Kolbert, Russ Tedrake, and Alberto Rodriguez. Parameter and contact force estimation of planar rigid-bodies undergoing frictional contact. *The International Journal of Robotics Research*, 2016.

[10] Nima Fazeli, Russ Tedrake, and Alberto Rodriguez. Identifiability analysis of planar rigid-body frictional contact. In *ISRR*, 2015.

[11] C. S. Gillmor. Coulomb and the evolution of physics and engineering in eighteenth-century france. *Princeton University Press*, 1971.

[12] Suresh Goyal, Andy Ruina, and Jim Papadopoulos. Planar sliding with dry friction part 1. limit surface and moment function. *Wear*, 143(2):307–330, 1991.

[13] Suresh Goyal, Andy Ruina, and Jim Papadopoulos. Planar sliding with dry friction part 2. dynamics of motion. *Wear*, 143(2):331–352, 1991.

[14] F. R. Hogan and J. R. Forbes. Trajectory tracking, estimation, and control of a pendulum-driven spherical robot. *Journal of Guidance*, 39(5):1118–1124, 2016.

[15] F. R. Hogan, J. R. Forbes, and T. D. Barfoot. Rolling stability of a power-generating tumbleweed rover. *Journal of Spacecraft and Rockets*, 51(67):1895–1906, 2014.

[16] François Hogan and Alberto Rodriguez. Feedback control of the pusher-slider system: A story of hybrid and underactuated contact dynamics. *Workshop on Algorithmic Foundation of Robotics (WAFR)*, 2016.

[17] R.D. Howe, I. Kao, and M.R. Cutkosky. Complementarity formulations and existence of solutions of dynamic multi-rigid-body contact problems with coulomb friction. *International Conference on Robotics and Automation (ICRA)*, 1988.

[18] Soo Hong Lee and M R Cutkosky. Fixture planning with friction. *Transactions of the ASME*, 113:320–327, 1991.

[19] K.M. Lynch, H. Maekawa, and K. Tanie. Manipulation and active sensing by pushing using tactile feedback. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1:416–421, 1992.

[20] M T Mason. Haptic perception, grasping and manipulation with simple robotic grippers, 2 2016. Remarks by Mathew T Mason at the CITRIS People and Robots Seminar Series, UC Berkeley [Accessed: 2017 04 20].

[21] Matthew T Mason. Compliance and force control for computer controlled manipulators. *IEEE Transactions on Systems, Man and Cybernetics*, 11(6):418–432, 1981.

[22] Matthew T Mason. Mechanics and planning of manipulator pushing operations. *The International Journal of Robotics Research*, 5(3):53–71, 1986.

[23] E. Meyer. Nanoscience: Friction and rheology on the nanometer scale. *World Scientific*, page 7, 2002.

[24] H. Moseley. On the equilibrium of the arch. *Trans. Cambridge Philosophical Soc.*, pages 293–314, 1835.

[25] J.S. Pang and J.C. Trinkle. Complementarity formulations and existence of solutions of dynamic multi-rigid-body contact problems with coulomb friction. *Mathematical Programming*, 73(2):199–226, 1996.

[26] Alberto Rodriguez and Matthew Mason. Two finger caging: Squeezing and stretching. *Workshop on the Algorithmic Foundations of Robotics(WAFR)*, 2008.

[27] D. Stewart and J.C. Trinkle. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *International Journal for Numerical Methods in Engineering*, 39(15):2673–2691, 1996.

[28] R Tedrake. Open problems session, 12 2016. Remarks by R Tedrake during the Open Problems Session of the 12th International Workshop on the Algorithmic Foundations of Robotics, San Francisco, California.

[29] C. Truesdell. Essays in the history of mechanics. *Springer-Verlag*, 1968.

[30] Wang Ya-Feng, Sun Fu-Chun, Zhang You-An, Liu Hua-Ping, and Min Haibo. Getting a suitable terminal cost and maximizing the terminal region for mpc. *Mathematical Problems in Engineering*, 2010.