## Conjugate Dynamic Programming

## Rodopoulos Charalampos

# Conjugate Dynamic Programming 

Master of Science Thesis

## For the degree of Master of Science in Systems and Control at Delft University of Technology

Rodopoulos Charalampos

November 25, 2021


## European

Research
Council
DCSC
Copyright © Delft Center for Systems and Control (DCSC) All rights reserved.

## Abstract

In decision making problems, the ability to compute the optimal solution can pose a serious challenge. Dynamic Programming (DP) aims to provide a framework to deal with a category of such problems, namely ones that involve sequential decision making. By dividing the original control problem into sub problems and solving it backwards in time, from the end of the time horizon to the start, the method can compute a map of optimal solutions with respect to the initial condition. In order to divide the original problem into subproblems the DP method takes advantage of the principle of optimality, which states that a sub-solution of the optimal solution should be the optimal solution for the equivalent subproblem. In control systems, where the state and decision spaces are continuous, the original DP framework can be intractable due to the size of the discretization needed to simulate the continuous space. Therefore, efficient approximations are needed to solve such problems. One promising method is called Conjugate Dynamic Programming (CDP). The CDP algorithm is able to transform the original framework and solve problems in the conjugate domain providing a computational advantage over the standard method. In this work, we aim to improve and extend the setting under which the CDP algorithm operates, thus providing a more concrete advantage over standard method . In that regard, we will extract the optimal control actions from within the CDP algorithm, eliminating the need for solving an extra optimization problem for their computation. In addition, we will introduce a different interpolation technique that can outperform the current one, in certain scenarios, thus granting the user more choice when solving a decision making problem.

## Table of Contents

Acknowledgements ..... ix
1 Introduction ..... 1
2 Theoretical Prerequisites ..... 3
2-1 Legendre-Fenchel transform and LLT algorithm ..... 3
2-2 Conjugate Dynamic Programming ..... 4
2-2-1 Original DP framework ..... 4
2-2-2 CDP algorithm ..... 5
3 Analytical Results ..... 7
3-1 Kuhn triangulation as an interpolation method ..... 7
3-1-1 Method outline and example in 2D ..... 7
3-1-2 Convex extension of a convex extensible function ..... 10
3-2 Extraction of optimal control input ..... 13
3-2-1 Theoretical Framework ..... 13
3-2-2 Example of SISO system ..... 14
4 Simulation Results ..... 17
4-1 Kuhn triangulation versus LERP ..... 17
4-2 Extraction of optimal control input ..... 20
4-2-1 Effects of discretization on extraction of optimal pairs ..... 20
4-2-2 1D example ..... 22
4-2-3 2D example ..... 28
5 Conclusions and Future Research ..... 33
5-1 Conclusions ..... 33
5-2 Future Research ..... 34
Master of Science Thesis Rodopoulos Charalampos
A CDP Algorithms ..... 35
Bibliography ..... 37

## List of Figures

2-1 Settings of CDP algorithm [1] ..... 6
3-1 Random query point and closest grid hypercube around it. ..... 8
3-2 Translated coordinate system so that the hypercube around the query point is the unit hypercube ..... 9
3-3 Plot of continuous function in 3D space. ..... 11
3-4 Scatter plot from data interpolated with Kuhn's triangulation. ..... 12
4-1 Error between cost-to-go functions calculated with interpolation methods and the benchmark. ..... 19
4-2 Optimal pairs for the two discretizations and the analytical solution (top) and percentage of error each pair for the two discretizations differs from the analytical solution (bottom). ..... 21
4-3 Percentage of error of conjugate function computed with LLT with respect to the analytically available conjugate function. ..... 22
4-4 Absolute error of one step ahead optimal control inputs against the benchmark averaged over 100 initial conditions. ..... 24
4-5 Number of invalid trajectories at each step in the horizon for the one step ahead predictions ..... 25
4-6 Average aggregate cost at each step for the control inputs calculated by the 3 different methods ..... 26
4-7 Trajectory error for each step averaged over 100 initial positions. ..... 27
4-8 Number of invalid trajectories at each step in the horizon. ..... 28
4-9 Number of invalid trajectories at each step in the horizon. ..... 29
4-10 Absolute error of one step ahead optimal control inputs against the benchmark averaged over 100 initial conditions. ..... 30
4-11 Number of invalid trajectories at each step in the horizon for the multistep imple- mentation ..... 31
4-12 Average aggregate cost at each step for the control inputs calculated by the 3 different methods. ..... 31
4-13 Trajectory error for each step averaged over 100 initial positions ..... 32

## List of Tables

4-1 Comparison between Kuhn triangulation and LERP as interpolation techniques. . 18

## Acknowledgements

The current thesis would not have been possible without my daily supervisors A. Sharifi Kolarijani and G.F. Max . Your feedback throughout this process was invaluable and allowed me to learn and grow as a person. Also a big thank you to my supervisor Dr. P. Mohajerin Esfahani for always being available when needed and for giving me the opportunity to work on such a challenging and rewarding problem.
To all my friends here in Delft, for all the good times we had together and for the fun inside and outside our studies. Also to my friends in Greece, who even though they were not here physically, they were always available when I needed someone to speak to.

## Chapter 1

## Introduction

The motivation behind this work lies on solving the Dynamic Programming (DP) problem by expanding on the work of [1]. Dynamic Programming aims to solve the Bellman equation,

$$
\begin{equation*}
J_{t}\left(x_{t}\right)=\min _{u_{t}}\left\{C_{t}\left(x_{t}, u_{t}\right)+J_{t+1}\left(x_{t+1}\right)\right\} \tag{1-1}
\end{equation*}
$$

by computing the optimal cost-to-go functions $J_{t}$ backwards in time, given the running cost $C_{t}\left(x_{t}, u_{t}\right)$ of choosing an action $u_{t}$ at a given state $x_{t}$, and the cost to go $J_{t+1}\left(x_{t+1}\right)$. at the next step $t+1$.

At the core of the DP formulation lies the principle of optimality introduced by Bellman [2]. By this we divide the original control problem into sub-problems which are simpler to solve. What makes the formulation of the DP problem attractive is that many interesting real world applications can be modeled and subsequently solved using DP. Some of these can be found in fields such as finance [3], power systems and the peak shaving problem $[4,5,6,7]$, as well as optimal control $[8,9]$, and inventory routing [10].

The ability to solve complex problems like the ones mentioned above in a computationally efficient manner marks the main point of this thesis. Approximating the cost-to-go function, also knows as value function, is the main interest behind solving the subclass of problems known as Approximate Dynamic Programming (ADP). The approximation is needed in problems where the computation of the exact cost-to-go function is impossible or requires an extremely large amount of resources . This can include problems in continuous spaces, where a sharp discretization would lead to the problem being intractable. An extensive overview of the theory and methods on how to solve such problems can be found in [11, 12, 13]. Popular methods include Q-Learning [14] and Temporal Difference (TD) Learning [15]. Furthermore, in [16], a method based on linear programming is proposed. The method tries to fit a linear combination of predetermined functions to approximate the cost-to-go function. Lastly, in $[17,18]$ a method is developed to compute max-plus linear approximations of the value function within a dictionary of functions. The necessity for ADP and generally more efficient ways to solve the DP problem stem from its biggest drawback, its high computational cost
when dealing with high dimensional state and action spaces, otherwise known as the curse of dimensionality [2]. In [1] an algorithm was developed, which exploits the Legendre-Fenchel transform [19],

$$
\begin{equation*}
F^{*}(s)=\sup _{x \in X}\{\langle x, s\rangle+F(x)\}, \quad \forall s \in S \subseteq \mathbb{R}^{N} \tag{1-2}
\end{equation*}
$$

in order to solve the DP equation (1-1). In equation (1-2), $F$ denotes a function with domain $X \subseteq \mathbb{R}^{n}$ and codomain $\mathbb{R}$, while $F^{*}$ denotes its conjugate function with domain $S \subseteq \mathbb{R}^{n}$ and codomain $\mathbb{R}$. In equation 1-2 $\langle x, s\rangle$ denotes the inner product between the two variables $s$ and $x$. In that regard, a linear time algorithm for computing the discrete Legendre transform [20], has been used for transforming the original problem and solving it in the conjugate domain. This has resulted in a more efficient implementation than solving the Bellman equation explicitly, which involves minimizing over the action space for any given state. Specifics on how the proposed algorithms in [1] work will be given in Chapter 2.
Contributions. The aim of this project is to improve the setting under which the LegendreFenchel transform can be used to solve the DP problem. In that regard, during this thesis we will aim to improve the computational aspect of working in the conjugate domain thus providing a more concrete advantage over traditional methods. More specifically, one key improvement revolves around extracting the optimal control action from within the CDP algorithm [1]. In the current implementation, in order to obtain the optimal control action, a minimization problem is solved during the forward stage of the algorithm, thus adding to the total computational cost. By leveraging the powers and capabilities of the LLT algorithm [20], we can extract the optimal control action during the backwards iteration at little to no extra computational cost, thus gaining an overall computational advantage. Furthermore, we aim to improve on other computational aspects of the algorithm by employing, possibly more efficient methods of interpolation/extrapolation without greatly affecting the overall computational cost of the algorithm.
Thesis organization. The current report is organized as follows. In Chapter 2, some theory that will be used throughout the thesis will be presented. This includes a short introduction into the CDP algorithm. In Chapter 3, we show theoretical results regarding the topics tackled in the present thesis. These include the theory behind extracting the optimal control input from within the CDP algorithm, and whether Kuhn triangulation can interpolate to a convex function, given that the discrete set of points are convex extensible. In Chapter 4, we present the numerical simulations and results associated with the theory of Chapter 3. In Chapter 5, we discuss the results and provide possible future research avenues.

## Chapter 2

## Theoretical Prerequisites

In this Chapter we will present some theoretical prerequisites that will be used throughout the thesis and are necessary for the reader to understand the following Chapters.

## 2-1 Legendre-Fenchel transform and LLT algorithm

The Legendre-Fenchel transform and the ability to compute it in linear time via the LLT algorithm [20] allows the CDP algorithm to outperform standard methods in computing the solution to DP problem. For completeness' sake, we briefly present the workings of the LLT algorithm. For more information regarding LLT the reader can refer to [20].

In equation (1-2), we presented the Legendre-Fenchel transform in continuous primal and dual spaces. The LLT algorithm computes the conjugation via an approximation, by discretizing the primal $X$ and dual $S$ spaces of equation (1-2). The discrete Legendre transform is presented in equation (2-1). The spaces $X_{d}, S_{d} \subset \mathbb{R}^{n}$ are both discrete.

$$
\begin{equation*}
F^{d *}(s)=\max _{x \in X_{d}}\left\{\langle x, s\rangle+F^{d}(x)\right\}, \quad \forall s \in S_{d} \tag{2-1}
\end{equation*}
$$

The superscript $d$, as in $F^{d}$, stands for a discrete function with a finite domain, while the subscript $d$ stands for a discrete (finite) set.
For the univariate case, in order to compute the conjugate function $F^{d *}(s)$, the algorithm takes the pairs of points in primal space,

$$
\begin{equation*}
\left\{\left(x_{1}, F^{d}\left(x_{1}\right)\right),\left(x_{2}, F^{d}\left(x_{2}\right)\right), \ldots,\left(x_{n}, F^{d}\left(x_{n}\right)\right)\right\} \tag{2-2}
\end{equation*}
$$

and computes a convex hull of the function $F^{d}$ if it is not convex. This can be done in linear time. Then the following slopes are computed:

$$
\begin{equation*}
c_{i}:=\frac{F^{d}\left(x_{i+1}\right)-F^{d}\left(x_{i}\right)}{x_{i+1}-x_{i}} . \tag{2-3}
\end{equation*}
$$

Since the function is convex these slopes are in ascending order. Given the points of interest in the dual space $S:=\left\{s_{1}, s_{2}, \ldots, s_{m}\right\}$, Lemma 3 of [20] proves that we can compute the conjugate function in the following manner:
(i) If $c_{i-1}<s_{i}<c_{i}, F^{d *}\left(s_{i}\right)=s_{i} x_{i}-F^{d}\left(x_{i}\right)$
(ii) If $c_{i}=s_{i}, F^{d *}\left(s_{i}\right)=s_{i} x_{i}-F^{d}\left(x_{i}\right)$ or $F^{d *}\left(s_{i}\right)=s_{i} x_{i+1}-F^{d}\left(x_{i+1}\right)$.

In the case we are dealing with multivariate functions, the discrete Legendre transform reads

$$
F^{d *}(s 1, s 2)=\max _{x_{1} \in X_{d 1}}\left[s_{1} x_{1}+\max _{x_{2} \in X_{d 2}}\left[s_{2} x_{2}-F^{d}\left(x_{1}, x_{2}\right)\right]\right], \quad \forall\left(s_{1}, s_{2}\right) \in S_{d 1} \times S_{d 2}
$$

with primal domain $X_{d}=X_{d 1} \times X_{d 2} \subset \mathbb{R}^{2}$ and dual domain $S_{d}=S_{d 1} \times S_{d 2} \subset \mathbb{R}^{2}$. This procedure is essentially a factorization which casts the original multivariate function as univariate ones and transforms these using the univariate method above. The algorithm for the two dimensional case will involve computing the intermediate function

$$
\begin{equation*}
F_{s_{2}}^{d}\left(x_{1}\right)=\max _{x_{2} \in X_{d 2}}\left[s_{2} x_{2}-F^{d}\left(x_{1}, x_{2}\right)\right], \quad \forall s_{2} \in S_{d 2}, \tag{2-4}
\end{equation*}
$$

and then computing the conjugate function

$$
\begin{equation*}
F^{d *}\left(s_{1}, s_{2}\right)=\max _{x_{1} \in X_{d 1}}\left[s_{1} x_{1}-\left(-F_{s_{2}}^{d}\left(x_{1}\right)\right)\right], \quad \forall s_{1} \in S_{d 1} \tag{2-5}
\end{equation*}
$$

This procedure can be generalised in higher dimensions.

## 2-2 Conjugate Dynamic Programming

## 2-2-1 Original DP framework

In [1], two novel algorithms are developed, in order to solve the DP in the conjugate domain. Below we present the transformation of the original DP problem to its equivalent in the conjugate domain and a short introduction into the two settings in the paper. We start with the discrete system:

$$
\begin{equation*}
x_{t+1}=f\left(x_{t}, u_{t}\right), \quad t=0,1, \ldots, T-1, \tag{2-6}
\end{equation*}
$$

where $f: \mathbb{R}^{\mathrm{n}} \times \mathbb{R}^{\mathrm{m}} \rightarrow \mathbb{R}^{\mathrm{n}}$ denotes the dynamics. The system is constrained in the input and state as follows:

$$
\begin{array}{ll}
x_{t} \in X \subset \mathbb{R}^{n}, & t \in\{0,1, \ldots, T\}, \\
u_{t} \in U \subset \mathbb{R}^{m}, & t \in\{0,1, \ldots, T-1\} . \tag{2-8}
\end{array}
$$

The stage and terminal cost functions are $C_{t}: X \times U \rightarrow \overline{\mathbb{R}}$ and $C_{T}: X \rightarrow \mathbb{R}$, where $\overline{\mathbb{R}}$ represents the extension of the real numbers to include $+\infty$. The reason for that is that if the state constraint is violated then the cost for that stage becomes $+\infty$.

The DP formulation for this problem reads:

$$
\begin{equation*}
J_{t}\left(x_{t}\right)=\min _{u_{t}}\left\{C\left(x_{t}, u_{t}\right)+J_{t+1}\left(x_{t+1}\right):(2-6)-(2-8)\right\}, \quad x_{t} \in X, \tag{2-9}
\end{equation*}
$$

Solving (2-9) backwards in time, we obtain the optimal cost-to-go functions as well as the control laws that will result in those functions. In order to solve such a problem, we need to iterate over all $x \in X$ and $u \in U$. When dealing with continuous spaces in $X$ and $U$ this procedure is intractable. Thus we need to approximate the solution by discretizing the state and input spaces. The discrete DP (d-DP) formulation reads:

$$
\begin{equation*}
J_{t}^{d}\left(x_{t}\right)=\min _{u_{t}}\left\{C\left(x_{t}, u_{t}\right)+\bar{J}_{t+1}^{d}\left(x_{t+1}\right)\right\}, \quad x_{t} \in X_{d} \subset X, \quad u_{t} \in U_{d} \subset U . \tag{2-10}
\end{equation*}
$$

The function $\bar{J}_{t+1}^{d}$ denotes the extension of the discretized cost-to-go function $J_{t+1}^{d}$. Since we are dealing with discrete spaces and only have access to the function values in a finite set of points the extension is needed to compute points outside of the given discrete domain.

For the remainder of this thesis every time we mention d-DP, we are referring to a brute force solution of the DP problem over the input and state spaces. In particular, by having finite discrete spaces, we can iterate over all $x \in X^{d}$ and $u \in U^{d}$ in order to find the optimal solution. The discrete spaces employed throughout this thesis are all grid like. This is done in order to take advantage of the speed up in computing the conjugate function using the LLT algorithm (see Remark 5 in [20]).

## 2-2-2 CDP algorithm

From the original DP formulation, we can create an equivalent problem in the dual domain. In [1], two different settings are presented for solving the DP problem in the conjugate domain. Below we will present the requirements that need to be met so that the appropriate setting can be used. For setting 1 of the CDP algorithm we need:
(i) Input-affine dynamics, $f(x, u)=f_{s}(x)+f_{i}(x) u$.
(ii) The sets $X$ and $U$ are compact and convex. Additionally, for every $x \in X$, the set of admissible inputs is non empty.
(iii) The stage cost is jointly convex in the state and input variable with a compact effective domain. The terminal cost needs to be also convex.

For the second setting we need the requirements described in setting 1 as well as:
(i) The input dynamics $f_{i}$ need to be state-independent. That means $f_{i}(\cdot)=B \in \mathbb{R}^{n \times m}$.
(ii) The stage cost is separable in the state and input variables. That means $C_{t}(x, u)=$ $C_{s}(x)+C_{i}(u), \quad C_{s}: X \rightarrow \mathbb{R} \quad$ and $C_{i}: U \rightarrow \mathbb{R}$.

If a problem meets the above requirements then we can transform the original DP. In Figure $2-1$, we present the alternate route of solving the DP problem for the two settings of the CDP algorithm developed in [1].


Figure 2-1: Settings of CDP algorithm [1]
The current Settings of the CDP algorithm compute the cost-to-go functions backwards in time via the alternate path presented in the figure. As it stands, none of the alternative paths compute the optimal control laws for a given step of the control problem. In order to acquire these laws the algorithm solves an additional minimization problem, for each step in the horizon, during the forward phase of the DP problem. The minimization problem is the following:

$$
u_{t}^{*}=\underset{u_{t} \in U_{d}}{\arg \min }\left\{C\left(x_{t}, u_{t}\right)+\bar{J}_{t+1}^{d}\left(x_{t+1}\right)\right\} .
$$

One of the contributions of the current thesis involves extracting the optimal control laws from the alternate path of the CDP algorithm and thus eliminating the minimization problems of presented above.

## Chapter 3

## Analytical Results

In this Chapter, we will be presenting the preliminary results regarding the problems tackled in the present thesis. This chapter is structured as follows: firstly, we will explore whether Kuhn's triangulation can create a convex extension out of a convex extensible set of discrete points. The reason why preserving convexity is important, is that the CDP algorithm is essentially blind to non convexity. This means that if the cost-to-go function $J_{t}$ is non convex, we will be introducing an error by approximating the function by its convex envelope. In the case of a convex cost-to-go function this error disappears. In addition, we will present the proof regarding the extraction of the optimal control input from within the CDP algorithm and present the calculations followed in a one dimensional example.

## 3-1 Kuhn triangulation as an interpolation method

## 3-1-1 Method outline and example in 2D

Using Kuhn's triangulation as an interpolation technique was originally presented in [21, 22]. Below we will present an outline of how the method works for completeness. This method is based on splitting the unit hypercube around the query point into simplexes. Then, by finding the simplex that contains our query point, we use a linear combination of the vertices of that simplex. The outline of the method is:

- First we need to translate and scale our coordinate system so that the cube around our query point becomes the unit hypercube. The new coordinates of the query point are $\left(x_{1}^{\prime}, \ldots, x_{d}^{\prime}\right)$, where $d$ is the number of dimensions of the space.
- Next we must employ a sorting algorithm in order to sort the coordinates of the translated query point. This will be used to identify the simplex in which the point lies.
- From the identified simplex, we use the coordinates of its vertices in order to create linear combinations of the query point's coordinates.
- After finding the coefficients from the previous step, we can use them as weights for the sum of the function values at each vertex.

In Figure 3-1, we present the random query point $\left(x_{1 q}, x_{2 q}\right)$ and the hypercube around it, which in two dimensions is a square.


Figure 3-1: Random query point and closest grid hypercube around it.
The first step involves translating and scaling the coordinate system in order for the hypercube around the query point to be the unit hypercube. For any point inside or on the perimeter of the hypercube we can use the following translation.

$$
\begin{gather*}
x_{1 p \_n e w}=F\left(x_{1 p}\right)=\left(x_{1 p}-x_{11}\right) /\left(x_{12}-x_{11}\right), \quad x_{11} \leq x_{1 p} \leq x_{12}  \tag{3-1}\\
x_{2 p \_n e w}=G\left(x_{2 p}\right)=\left(x_{2 p}-x_{21}\right) /\left(x_{22}-x_{21}\right), \quad x_{21} \leq x_{2 p} \leq x_{22} \tag{3-2}
\end{gather*}
$$

where $p$ denotes any point inside or on the perimeter of the hypercube. The results of the above procedure lead to the new coordinate system of Figure 3-2.
Using equations (3-1)-(3-2) the new coordinates of our query point become:

$$
\begin{gather*}
x_{1 q \_ \text {new }}=\left(x_{1 q}-x_{11}\right) /\left(x_{12}-x_{11}\right)  \tag{3-3}\\
x_{2 q \_ \text {new }}=\left(x_{2 q}-x_{21}\right) /\left(x_{22}-x_{21}\right) . \tag{3-4}
\end{gather*}
$$

Next, we will divide the hypercube in $d!$ simplexes. The coordinates of the points inside each simplex have to satisfy the following equation.

$$
0 \leq x_{r(1)} \leq x_{r(2)} \leq \ldots \leq x_{r(d)} \leq 1
$$



Figure 3-2: Translated coordinate system so that the hypercube around the query point is the unit hypercube
where $r$ represents a possible permutation of $(1,2,3, \ldots, d)$.
In our example, we will separate the square in triangles (2D simplexes) by drawing the diagonal of the square connecting the origin $(0,0)$ to the following vertex $(1,1)$. In two dimensions this procedure creates two triangles which can be seen in Figure 3-2. The lower triangle in the figure is described by the equation $x_{2} \leq x_{1}$, while the upper triangle is described by the equation $x_{1} \leq x_{2}$. Next, the procedure requires that we use a sorting algorithm to sort the coordinates of the new query point. The sorting procedure is the most important step of the algorithm. It will allow us to easily identify the simplex inside which the point lies and its corresponding vertices. We know that every simplex inside the hypercube will have two known vertices in $(0,0,0, \ldots)$ and ( $1,1,1, \ldots)$. Thus we only need to identify $d-1$ points, where $d$ is the dimension we are working with. An easy way to identify the simplex requires using the dimension on which the sorted coordinates lie. We start with a matrix of size $(d+1) \times d$ filled with zeros. Then starting from the largest coordinate we fill the appropriate column with 1's for d rows. Then we move onto the next coordinate where we fill the appropriate column with 1 for d-1 rows. The procedure continues to the smallest coordinate where we fill the appropriate column with 1 for one row. Every row of the matrix will represent the coordinates of the vertices of each simplex. The procedure for the 2D case can be seen below. We start from the matrix filled with 0 :

$$
\left[\begin{array}{ll}
0 & 0  \tag{3-5}\\
0 & 0 \\
0 & 0
\end{array}\right]
$$

Since in our example the coordinate of the first dimension is larger that that of the second we will fill the first column with 1s for the first two rows. This leads to

$$
\left[\begin{array}{ll}
1 & 0  \tag{3-6}\\
1 & 0 \\
0 & 0
\end{array}\right] .
$$

Then we move on to the second largest dimension, and last in our example, $x_{2}$. We will fill the second column of the matrix with 1 s for one row. The final result can be seen below

$$
\left[\begin{array}{ll}
1 & 1  \tag{3-7}\\
1 & 0 \\
0 & 0
\end{array}\right] .
$$

If the query point lied in the upper triangle, the coordinate of the second dimension would have been bigger and following the same procedure we would first fill the second column with 1 s for two rows and then the first column with 1 for one row. Every row of that matrix represents a vertex of the simplex the query point lies inside of. Comparing the matrix in equation (3-7) and the Figure 3-2 we can see that the coordinates of the vertices of the simplex have been identified correctly.

Having acquired the vertices of the simplex, we can now express the coordinates of the query point as a linear combination of the coordinates of each vertex. By solving the following system of equations.

$$
\left[\begin{array}{lll}
1 & 1 & 1  \tag{3-8}\\
0 & 1 & 1 \\
0 & 0 & 1
\end{array}\right]\left[\begin{array}{l}
a_{1} \\
a_{2} \\
a_{3}
\end{array}\right]=\left[\begin{array}{c}
1 \\
x_{1 \text { qnew }} \\
x_{2 \text { qnew }}
\end{array}\right]
$$

Upon solving the preceding equation we obtain the coefficients $a_{1}, a_{2}, a_{3}$. Then, we can now compute the function value at the query point as follows:

$$
\begin{equation*}
F\left(x_{1 q}, x_{2 q}\right)=a_{1} F\left(x_{11}, x_{21}\right)+a_{2} F\left(x_{12}, x_{21}\right)+a_{3} F\left(x_{12}, x_{22}\right) \tag{3-9}
\end{equation*}
$$

Essentially the function value at the query point is calculated as a linear combination of the function values at the vertices of the simplex, based on the proximity of the query point to the vertices themselves.

## 3-1-2 Convex extension of a convex extensible function

Definition 3-1.1. A discrete function $F^{d}: X_{d} \rightarrow \overline{\mathbb{R}}$ is said to be convex extensible when there exists a continuous convex function $F: \mathbb{R}^{n} \rightarrow \overline{\mathbb{R}}$ where $\forall x \in X_{d}$ the following holds $F(x)=F^{d}(x)$.

The CDP algorithms in settings 1 and 2 both work with discrete functions evaluated at a grid-like discrete set of points. At certain steps, it is necessary to evaluate these functions at points outside of the available points. Thus, we need a interpolation technique. Assuming we start from a convex extensible set of points, we would like this technique to identify a continuous convex function that satisfies the definition of convex extensibility listed above.

The theoretical framework that aims to describe the extension of convexity properties in discrete functions is called Discrete Convex Analysis [23, 24, 25].

Unfortunately, convex extensibility is not a sufficient condition for Kuhn triangulation to lead to a convex extension for a discrete function. We will show that by presenting the following example. Assume the continuous function

$$
\begin{equation*}
f\left(x_{1}, x_{2}\right)=\max \left(x_{1}-3 x_{2},-2 x_{1}+3 x_{2}\right) . \tag{3-10}
\end{equation*}
$$

The above function is convex as the maximum of two affine functions. In Figure 3-3, we present the function plotted in 3D space.


Figure 3-3: Plot of continuous function in 3D space.
We now assume the following discretization of the domain of the function. The function value at these points only, is assumed to be known. The discrete function is as follows:

$$
\begin{equation*}
f^{d}\left(x_{1}, x_{2}\right)=\max \left(x_{1}-3 x_{2},-2 x_{1}+3 x_{2}\right), \quad x_{1}, x_{2} \in X_{d}^{1}=[0,2,4,6, \ldots, 20]^{2} . \tag{3-11}
\end{equation*}
$$

We will be using Kuhn triangulation to interpolate the function value at the point $x_{q}=$ $\left(x_{1}, x_{2}\right)=(2,1)$. The continuous function value at this point is $f(2,1)=-1$. The points around our query point are $x_{l l}=[0,0], x_{l r}=[2,0], x_{u l}=[0,2], x_{u r}=[2,2]$. We transform these points so that $x_{l l}=[0,0]$ and $x_{u r}=[2,2]$ become the opposite corners of the unit square. Now our points are: $x_{l l}=[0,0], x_{l r}=[1,0], x_{u l}=[0,1], x_{u r}=[1,1]$. Our query point $x_{q}$ lies in the 2D simplex created by the points $x_{l l}=[0,0], x_{l r}=[1,0]$ and $x_{u r}=[1,1]$. Our point $x_{q}$ is the following linear combination of the simplex's corners:

$$
x_{q}=\left[\begin{array}{lll}
a_{1} & a_{2} & a_{3}
\end{array}\right]\left[\begin{array}{c}
x_{l l}  \tag{3-12}\\
x_{l r} \\
x_{u r}
\end{array}\right]=\left[\begin{array}{lll}
0 & 0.5 & 0.5
\end{array}\right]\left[\begin{array}{c}
x_{l l} \\
x_{l r} \\
x_{u r}
\end{array}\right]
$$



Figure 3-4: Scatter plot from data interpolated with Kuhn's triangulation.

We then use these coefficients to calculate the value of the function at our query point:

$$
\begin{gathered}
\overline{f^{d}}\left(x_{q}\right)=\left[\begin{array}{lll}
0 & 0.5 & 0.5
\end{array}\right]\left[\begin{array}{c}
f^{d}\left(x_{l l}\right) \\
f^{d}\left(x_{l r}\right) \\
f^{d}\left(x_{u r}\right)
\end{array}\right]=0 \cdot f\left(x_{l l}\right)+0.5 \cdot f\left(x_{l r}\right)+0.5 \cdot f\left(x_{u r}\right) \Rightarrow \\
\bar{f}^{d}\left(x_{q}\right)=0 \cdot 0+0.5 \cdot 2+0.5 \cdot 2=2 .
\end{gathered}
$$

The notation $\overline{f^{d}}$ indicates extending the discrete function to points outside its domain.

It is evident that the two values are not the same. Thus the method fails to recover the original convex function. Since there can be more than one convex function fitting a set of convex extensible points we need to investigate this possibility as well. If we use Kuhn triangulation to extend the discrete function of equation (3-11) to the following domain

$$
X_{d}^{2}=[0,1,2,3, \ldots, 20] \times[0,1,2,3, \ldots, 20]
$$

we obtain the result shown in Figure 3-4. The new domain $X_{d}^{2}$ consists of the original points in the domain $X_{d}^{1}$, for which we have the functions values available, and the points on which we used Kuhn triangulation to extend the original discrete function.

What we can observe in Figure 3-4 is that the extended discrete function presents this region where the function values form a non convex area. The method fails because the original discretization is not dense enough to capture the direction of the discretization in every square of the domain. Thus, in the example presented, the true function value of the query point lies below the function values of both points used to compute it. Although the method cannot guarantee the extended function will also be convex extensible, it is important to check whether the method provides a more accurate result when extending a discrete function than the method currently used in the CDP algorithm, namely multilinear interpolation (LERP).

## 3-2 Extraction of optimal control input

In [1], the proposed algorithms only return the cost-to-go functions but not the optimal control actions that will lead to those cost-to-go functions. For that reason, an additional minimization problem has to be solved in order to acquire the optimal control laws. In order to avoid the extra computational cost of the minimization problem, we will inquire whether optimal control laws can be extracted from within the CDP algorithm. This idea is based on the fact that the LLT algorithm [20] allows us to keep the optimal pairs between the primal and dual domains. This should allow us to find a relationship between the current state $x$ and the optimal input $u$ that will minimize the one step cost function.

## 3-2-1 Theoretical Framework

Lemma 3-2.1. For the Settings in the CDP algorithm we can extract the optimal control action by employing the optimal pairs in the conjugation steps as follows.

$$
\begin{array}{r}
y^{*}(x)=\underset{y}{\arg \max }\left\{\left\langle y, f_{s}(x)\right\rangle-\phi_{x}(y)\right\}, \\
u^{*}(y)=\underset{u}{\arg \max }\left[\left\langle-f_{i}^{T}(x) y^{*}, u\right\rangle-C(x, u)\right] . \tag{3-14}
\end{array}
$$

Proof. Following the proof of Appendix B.2.1 of [1], we start from the reformulated DP problem.

We know

$$
\begin{equation*}
\phi_{x}(y)=C_{x}^{*}\left(-f_{i}^{T}(x) y\right)+J^{*}(y), \tag{3-15}
\end{equation*}
$$

and

$$
\hat{\mathcal{T}}[J](x)=\phi^{*}\left(f_{s}(x)\right)=\max _{y} \max _{u}\left\{\left\langle y, f_{s}(x)\right\rangle-\left\langle-f_{i}^{T}(x) y, u\right\rangle+C(x, u)-J^{*}(y)\right\} .
$$

We define

$$
L_{1}(x, u, y):=\left\langle y, f_{s}(x)\right\rangle-\left\langle-f_{i}^{T}(x) y, u\right\rangle+C(x, u)-J^{*}(y) .
$$

Then,

$$
\begin{aligned}
L_{2}(x, y) & =L_{1}\left(x, u^{*}(y), y\right)=\left\langle y, f_{s}(x)\right\rangle-\left\langle-f_{i}^{T}(x) y, u^{*}\right\rangle+C\left(x, u^{*}\right)-J^{*}(y) \\
& \stackrel{(3-14)}{=}\left\langle y, f_{s}(x)\right\rangle-C_{x}^{*}\left(-f_{i}^{T}(x) y\right)-J^{*}(y) \\
& \stackrel{(3-15)}{=}\left\langle y, f_{s}(x)\right\rangle-\phi_{x}(y) .
\end{aligned}
$$

Finally,

$$
\begin{aligned}
L_{3}(x) & =L_{2}\left(x, y^{*}(x)\right)=\left\langle y^{*}, f_{s}(x)\right\rangle-\phi_{x}\left(y^{*}\right) \\
& \stackrel{(3-13)}{=} \phi^{*}\left(f_{s}(x)\right)
\end{aligned}
$$

We have proven that

$$
\hat{\mathcal{T}}[J](x)=L_{1}\left(x, u^{*}, y^{*}\right),
$$

using the mappings in equations (3-13)-(3-14)

The procedure to obtain the optimal control input from the current state is outlined below

1. From the current state $x$ compute $f_{s}(x)$.
2. Find the optimal dual variable $y$ from $f_{s}(x)$, using eq. (3-13).
3. Compute $-f_{i}^{T}(x) y$.
4. Find the optimal input variable $u$ from $-f_{i}^{T}(x) y$, using eq. (3-14).

## 3-2-2 Example of SISO system

Below we will present the procedure outlined above in a SISO system that can be solved analytically. The time horizon will be one step $(T=1)$. The procedure can be extended to longer than one step horizons. In Chapter 4, we will also present results from simulations in more complex systems.
We assume the following system

$$
\begin{array}{r}
x_{t+1}=2 x_{t}+u_{t} \\
C_{t}\left(x_{t}, u_{t}\right)=x_{t}^{2}+u_{t}^{2} \\
C_{T}=x_{T}^{2} \tag{3-18}
\end{array}
$$

For this system, we first compute the conjugate of the stage cost with respect to the input variable $u$ as follows.

$$
C^{*}(y)=\max _{u}\{\langle y, u\rangle-C(x, u)\}=\max _{u}\left\{y u-x^{2}-(u)^{2}\right\}=\frac{y^{2}}{4}-x^{2} .
$$

Next we compute the conjugate of the cost-to-go function at the terminal step which is the terminal cost $C_{T}$. Since we are only solving for one step, from here onwards, we will not use the subscript $t$ to indicate the step of the horizon.

$$
\begin{aligned}
J^{*}(y) & =C_{T}^{*}(y)=\max _{x}\left\{\langle y, x\rangle-C_{T}(x)\right\} \\
& =\max _{x}\left\{y x-x^{2}\right\}=\frac{y^{2}}{4}
\end{aligned}
$$

We can now form the function in $(3-15)$ and then compute the optimal pairs between the primal state and the dual state variables.

$$
\begin{align*}
y^{*}(x) & =\underset{y}{\arg \max }\left\{\left\langle f_{s}(x), y\right\rangle-C_{x}^{*}\left(-f_{i}^{T}(x) y\right)-J^{*}(y)\right\} \\
& =\underset{y}{\arg \max }\left\{2 x y-\frac{y^{2}}{4}+x^{2}-\frac{y^{2}}{4}\right\}  \tag{3-19}\\
& =2 x .
\end{align*}
$$

Next, we will compute the optimal pair between the input variable $u$ and the dual state variable $y$.

$$
\begin{align*}
u^{*}(y) & =\underset{u}{\arg \max }\left\{\left\langle-f_{i}^{T}(x) y, u\right\rangle-C(x, u)\right\} \\
& =\underset{u}{\arg \max }\left\{-y u-x^{2}-u^{2}\right\}  \tag{3-20}\\
& =-\frac{y}{2}
\end{align*}
$$

By combining the optimal pairs acquired in (3-19) and (3-20), we end up with the optimal control law

$$
\begin{equation*}
u^{*}(x)=u^{*}(y(x))=\frac{-y(x)}{2}=-x \tag{3-21}
\end{equation*}
$$

On the other hand, by solving the original DP formulation analytically we would get

$$
\begin{align*}
u^{*}(x) & =\underset{u}{\arg \min }\left\{C(x, u)+J\left(x_{T}\right)\right\} \\
& =\underset{u}{\arg \min }\left\{x^{2}+u^{2}+x_{T}^{2}\right\} \\
& =\underset{u}{\arg \min }\left\{x^{2}+u^{2}+(2 x+u)^{2}\right\}  \tag{3-22}\\
& =\underset{u}{\arg \min }\left\{x^{2}+u^{2}+4 x^{2}+u^{2}+4 x u\right\} \\
& =\underset{u}{\arg \min }\left\{5 x^{2}+2 u^{2}+4 x u\right\} \\
& =-x
\end{align*}
$$

As we can see, the optimal control law obtained in both cases is exactly the same. A few further remarks regarding the proposed way of extracting the optimal control input. The analytical procedure outlined above is being computed with all spaces being continuous. That means that whatever the optimal pairings are, we know that the space contains them and we can extract them without error. The conjugations in the CDP algorithm are performed over discrete spaces with the LLT algorithm. As input to the LLT algorithm, we need to provide both the primal and dual grids. Since we have no way of knowing the optimal pairing between the two grids beforehand, the choice of the two grids will affect the result of the conjugation. The difference this will have in the real world simulations will be examined in Chapter 4.

## Simulation Results

In this chapter, we will present the simulation results regarding the various topics tackled in this thesis. Firstly, we will present the difference between using Kuhn triangulation as an interpolation technique and multilinear interpolation (LERP). Then, we will use the theory presented in the previous chapter to extract the optimal control input from within the CDP algorithm.

## 4-1 Kuhn triangulation versus LERP

In this section, we will present the results of using Kuhn triangulation versus LERP as an interpolation technique. The interpolation techniques will be tested on line 6 of Setting 2 of the CDP algorithm in order to calculate $\bar{\phi}^{d * d}[1, \mathrm{p} .19]$. For the reader's convenience the algorithms of [1] are included in the appendix of the current thesis. For both techniques we will calculate the optimal cost-to-go functions $J_{t}$, that will be used to extract the optimal control inputs via a minimization. We will calculate the cost-to-go function for a horizon of $T=10$ for the following system.

$$
\begin{array}{r}
x_{t+1}=\left[\begin{array}{cc}
-0.5 & 2 \\
1 & 3
\end{array}\right] x_{t}+\left[\begin{array}{cc}
1 & 0.5 \\
1 & 1
\end{array}\right] u_{t} \\
C_{t}\left(x_{t}, u_{t}\right)=\left\|x_{t}\right\|_{2}^{2}+e^{\left|u_{t 1}\right|}+e^{\left|u_{t 2}\right|}-2 \\
C_{T}=\left\|x_{T}\right\|_{2}^{2} \tag{4-3}
\end{array}
$$

The system is constrained as follows

$$
\begin{align*}
& x_{t} \in[-1,1]^{2}  \tag{4-4}\\
& u_{t} \in[-2,2]^{2} \tag{4-5}
\end{align*}
$$

For our testing, the system is discretized in 11, 21 and 41 points in each dimension of the state and input grids. In Table 4-1, we will present the performance of the two interpolation techniques with respect to the overall cost of the trajectory that was generated from the control inputs obtained from the cost-to-go functions. These control actions were generated for a total of 100 initial conditions. In more detail, we present the interpolation technique used, the discretization of the state and input grid, the percentage of cases each interpolation technique outperforms the other, and the average cost improvement in those cases.

Table 4-1: Comparison between Kuhn triangulation and LERP as interpolation techniques.

| Method \& Discretization | C.I w.r.t other method (\% of cases) ${ }^{1}$ | Avg Cost Improvement |
| :---: | :---: | :---: |
| Kuhn triangulation - 11 | 17 | 0.5080 |
| LERP - 11 | 17 | 0.7140 |
| Kuhn triangulation - 21 | 15 | 0.3097 |
| LERP - 21 | 17 | 0.2959 |
| Kuhn triangulation - 41 | 11 | 0.0840 |
| LERP -41 | 15 | 0.1461 |

As we can see each, interpolation technique outperforms the other in about $15 \%$ of the cases and that number does not change considerably when increasing the discretization. Furthermore, the average cost improvement ranges from $0.08-0.71$ which translates to a percentage improvement of $1.5-8 \%$. This improvement, although important in the sparser grids, does not translate to any algorithm being the superior over the other. In about two thirds of the cases the algorithms performed exactly the same, which can be attributed to the sparse discretisation of the input grid in the sparser cases. The two methods could be computing different values for the control input but since the discretization is not dense enough, the difference is lost in the end result. In the more dense cases, the dense descritization of the state grid is leading to smaller errors and thus the methods perform similarly. Another reason why the two algorithms perform similarly is that when the query point, we are trying to compute the function value for, lies on or very close to the grid lines, then Kuhn's triangulation becomes the same as linear interpolation.

For the sake of completeness in Figure 4-1, we present the error of the cost-to-go functions obtained from the two interpolation methods versus our benchmark cost-to-go function for different discretizations for the first step in the horizon. The benchmark cost-to-go function was computed using a high density d-DP solution. As is evident from the figure, both methods perform very similarly to each other, which supports the comparison presented in Table 4-1. Also, as the discretization becomes more dense, the methods approach the benchmark which is something to be expected.

Overall, it is evident that the two methods perform very close to one another, and any small advantage one has over the other could be due to the different cost-to-go functions and not a clear advantage of one method.

[^0]

Figure 4-1: Error between cost-to-go functions calculated with interpolation methods and the benchmark.

## 4-2 Extraction of optimal control input

As we mentioned in Chapter 3, although theoretically the extraction of the optimal control input can be achieved, there are factors that could influence the result like the choice of the dual state grid. In Chapter 3, we extracted the optimal pairs analytically in the continuous domain. Since the CDP algorithm is working in the discrete domain and since we need to specify both the primal and dual grids before the conjugation, it is likely that the theoretical optimal pairs are not contained in the specified grids. In that case, the LLT algorithm will return a set of sub-optimal optimizers between the primal and dual grid. These pairs will then introduce an error in the control action we will compute. In the current version of the CDP algorithm the control inputs are computed by solving a minimization problem during the forward iteration of the DP problem. In the next sections, we will compare these with the control inputs obtained from our proposed method.

## 4-2-1 Effects of discretization on extraction of optimal pairs

As we mentioned already, extracting the optimal pairs between the primal and dual domain during the conjugation could be affected by the chosen grids. For our example, we will use the following function to perform the conjugation on.

$$
\begin{equation*}
f\left(x_{1}, x_{2}\right)=x_{1}^{2}+x_{2}^{2}, \quad x_{1}, x_{2} \in \mathbb{R} \tag{4-6}
\end{equation*}
$$

for which we can analytically compute the conjugate function

$$
\begin{equation*}
f^{*}\left(s_{1}, s_{2}\right)=\frac{s_{1}^{2}}{4}+\frac{s_{2}^{2}}{4} \tag{4-7}
\end{equation*}
$$

with the optimizer between the primal and dual domains being

$$
\begin{align*}
& x_{1}^{*}\left(s_{1}\right)=\underset{x_{1}}{\arg \max }\left\{\left\langle x_{1}, s_{1}\right\rangle+f_{x_{2}}\left(x_{1}\right)\right\}=\frac{s_{1}}{2},  \tag{4-8}\\
& x_{2}^{*}\left(s_{2}\right)=\underset{x_{2}}{\arg \max }\left\{\left\langle x_{2}, s_{2}\right\rangle+f_{x_{1}}\left(x_{2}\right)\right\}=\frac{s_{2}}{2} . \tag{4-9}
\end{align*}
$$

Now consider the following discretization of the function above.

$$
\begin{array}{r}
f^{d}\left(x_{1}, x_{2}\right)=x_{1}^{2}+x_{2}^{2} \\
\left(x_{1}, x_{2}\right) \in\{-4,-3.6,-3.2,-2.8, \ldots, 2.8,3.2,3.6,4\}^{2} \tag{4-11}
\end{array}
$$

In order to perform the conjugation via the LLT, we will choose two different grids for the dual domain.

$$
\begin{array}{r}
\left(s_{1}, s_{2}\right) \in S_{\text {opt }}=\{-8,-7.2,-6.4, \ldots, 6.4,7.2,8\}^{2} \\
\left(s_{1}, s_{2}\right) \in S_{\text {subopt }}=\{-11,-9.9,-8.8, \ldots, 8.8,9.9,11\}^{2} \tag{4-13}
\end{array}
$$



Figure 4-2: Optimal pairs for the two discretizations and the analytical solution (top) and percentage of error each pair for the two discretizations differs from the analytical solution (bottom).

All the chosen grids contain 21 equidistant points in every dimension. The reason we chose the two above discretizations of the dual domain is that the first one is the optimal discretization based on the analytical result in equations (4-8) and (4-9), while the second contains a larger space that given the same number of points should not contain all the optimal pairs. In Figure $4-2$, we plot the pairs obtained from each discretization on top of the optimizer mappings we obtained in equations (4-8)-(4-9), as well as the percentage each pair differs from its analytically calculated optimal. Given that we calculate the error as a percentage of the optimal value, we must note that in the case of the pair being $(0,0)$, both times the algorithm has identified the pair correctly and we set the error manually to zero. As we can observe from the figure, when the dual grid is chosen to be exactly the optimal of the primal grid, LLT computes the optimal pairs without error. In the case of the sub-optimal grid we can see that the error in the sub-optimal pairs can reach $30 \%$ of the optimal value. In our case this can lead to sub-optimal control inputs that will increase the overall cost for a trajectory. The LLT algorithm, in higher than one dimensions, employs a factorization scheme (see Section 2.1). This means that in order to extract the control input in the second dimension, we need the pair from the first dimension. This could introduce an additional error in the control inputs of the dimensions following the first. A higher discretization of the space could potentially lead to better results, since a denser grid is more likely to contain the points needed for the optimal pairs. Given that we are using the LLT algorithm to solve a DP problem, we must consider the curse of dimensionality. A dense grid in high dimensions could lead to large time requirements for computing the solution, or even make the problem intractable. To summarize, the first important result that we should keep in mind is that the choice of the dual grid can greatly affect our ability to extract the optimal pairs.

In the CDP algorithm's current state, the optimal pairs of the conjugations are not employed. In every call to the LLT algorithm only the conjugate function is needed, in order to compute the cost-to-go functions. These functions are later used to extract the control inputs via a minimization problem. It is important to see how a sub-optimal discretization can affect the values of the conjugate function, since the function is used as an alternative to extract the optimal control inputs. In Figure 4-3, we plot how much the conjugate functions differ from the analytically available result in equation (4-7) for the two discretization mentioned above.


Figure 4-3: Percentage of error of conjugate function computed with LLT with respect to the analytically available conjugate function.

Again, as these Figures represent the percentage of error, in the case of $(0,0)$, the function value was correctly identified and thus the error was set to zero.
The function values when the discretization is optimal, match the ones from the analytical result. For the sub-optimal discretization we can see that the function values can be off by up to $8 \%$. This result means, that in this case, and for the same sub-optimal choice for the dual grid, the conjugate function is closer to the analytically available conjugate, than the conjugation optimizers are to the analytically available optimizer. Since the current method of extracting the control inputs in the CDP algorithm, employs only the conjugate functions, it could be that it will lead to a smaller error compared to our proposed method which employs the conjugate pairs.

This will be tested also, in the following section on two different systems, a 1D and a 2D.

## 4-2-2 1D example

Firstly, we will apply the procedure to the model described below. The dynamics is given by:

$$
x_{t+1}=2 x_{t}+u_{t},
$$

The system is constrained as follows:

$$
\begin{align*}
& -1 \leq x_{t} \leq 1  \tag{4-14}\\
& -2 \leq u_{t} \leq 2 \tag{4-15}
\end{align*}
$$

The stage and terminal costs are:

$$
\begin{array}{r}
C_{t}=x_{t}^{2}+e^{\left|u_{t}\right|}, \\
 \tag{4-17}\\
\quad C_{T}=x_{T}^{2} .
\end{array}
$$

We will solve the DP problem and obtain the optimal control inputs via 3 different methods for a time horizon of $T=10$. First is the procedure using the optimal conjugation pairs described in the previous Chapter. Then the procedure that the CDP currently uses, which is the minimization of the cost-to-go function during the forward iteration of the DP problem. Lastly, for our benchmark we will use the d-DP solution from a high density discretized grid (81 points in each dimension). .

In Figure 4-4, we present for each step in the horizon the absolute error between the control input obtained from the two methods and the benchmark, averaged out over 100 initial conditions. For the sake of comparing the two methods fairly, at each step we start from the same initial condition. That means that the starting state in each step is the same for all methods. At each step we assume that $x_{t}$ is the same for both methods and we calculate $u^{*}\left(x_{t}\right)$ with the proposed method, the forward minimization and the d-DP solution. The areas filled with color depict the area where $80 \%$ of the samples lie. The initial conditions were selected randomly from a uniform distribution over the interval $[-1,1]$. The reason we used a uniform distribution is that we wanted to include as much of the state space as possible so as the performance of the methods is not influenced by parts of the state space where one method could outperform the other. The subfigures present how the error changes for different discretizations of the input and state space. The proposed method could not control all the initial conditions at each step, so in the figure we present only those conditions that were able to be controlled. In Figure $4-5$, we present the number of initial conditions that could not be controlled at each step. Note that we add the initial conditions that were not controlled at the previous steps to the current one.

As we can see, in the Figure 4-4 the methods perform extremely well. As we increase the discretization of the state and input spaces both errors are decreasing reaching very satisfactory levels. The difference between the errors seems to be decreasing with a higher discretization, which should be expected with more dense grids. Since we start from the same initial conditions spanning the entire state space $x$ we can see that both methods perform similarly at each step. This is expected as for the first few steps the cost-to-go functions are the same for each method, which should in turn give us the same results for the control inputs.

Since we are working with a time horizon of $T=10$, it is important to check how the methods perform over the entire horizon and how sub-optimal control actions computed by each method can affect the next steps of the horizon. For that reason, in Figure 4-6, we present the cost-to-go at every step when using the control actions we obtained above from each method and the benchmark. The costs are again averaged over the 100 initial conditions and at every step the cost is accumulated over the previous time steps. The filled areas represent the area inside which $80 \%$ of the our costs for the current step lie. The proposed method was not able to control all initial conditions, with a number of trajectories breaking the constraints. For that reason in the Figures $4-6$ and $4-7$ we have only included the trajectories that did not violate the constraints. In Figure 4-8, we present the number of invalid trajectories, for the proposed method, at each step of the horizon for different discretizations. Note that at each step we accumulate the number of trajectories over time. As we can see in the sparse discretization the proposed method cannot control all trajectories. The reason being that the extraction of the optimal conjugation pairs with that discretization of the state and input spaces, is not possible, leading to higher errors. Both the forward minimization method and the benchmark were able to control all trajectories without violating the constraints.


Figure 4-4: Absolute error of one step ahead optimal control inputs against the benchmark averaged over 100 initial conditions.


Figure 4-5: Number of invalid trajectories at each step in the horizon for the one step ahead predictions.

In Figure 4-6, the forward minimization method outperforms the conjugation pairs by a considerable amount. The difference between the two decreases as the grid becomes more dense. The difference between the two methods can be attributed to a snowballing effect of the errors at each step. As the proposed method computes slighty worse control inputs at each step, new sub-optimal states will be reached by the system. These sub-optimal states will then lead to larger inputs in order to be controlled and thus a propagation effect is created where each choice of control input affects the cost at later stages, since its effect needs to be mitigated. That is the reason why the cost in the case of the conjugate pairs keeps increasing, or takes longer to reach a steady state than the other two cases.

Lastly, we present in Figure 4-7 the absolute error of the trajectory of the system when controlled by the actions obtained by each method versus the benchmark averaged over the 100 initial conditions. Even though the snowballing effect described above will heavily influence the following figure, we think its important to include it since these are the "optimal" trajectories the methods were able to compute, for the given horizon. The error is indicative of the methods' multi-step performance and thus an important result in the comparison between the two methods.

In the figure, we observe a similar relation to the previous plots, where the forward minimization control inputs outperform the control inputs from the conjugate pairs, and the difference becomes smaller as we increase the discretization. There is one more interesting aspect we need to discuss in these plots. In the last step of the horizon, the error of the trajectories is increasing. This could be explained because as we get closer to the horizon and the state remains within the constraints, it is no longer optimal to incur large control inputs on the system, as it would lead to higher costs overall. So the system is left to drift, in a sense. Given the dynamics being unstable, without control the sub-optimal state computed by the two methods will drift faster than the optimal state computed by the benchmark. This lead to the increase of the overall error. This can also be backed up by the results in Figure 4-8 where only at the last step we have uncontrollable states.


Figure 4-6: Average aggregate cost at each step for the control inputs calculated by the 3 different methods.


Figure 4-7: Trajectory error for each step averaged over 100 initial positions.


Figure 4-8: Number of invalid trajectories at each step in the horizon.

## 4-2-3 2D example

In this section, we will be presenting the same plots for a 2 -state, 2 -input system. The system is described by the Equations (4-1)-(4-5).

We will be solving for a finite time horizon of $T=10$. Again, the control actions are computed with the the proposed method, the current method implemented in the CDP algorithm and our benchmark, which is a high density d-DP solution. We will test the performance of the methods against each other and the benchmark. The results are averaged over 100 initial conditions. The proposed method was not able to control the system for each initial condition. In Figure 4-9, we present the accumulating number of trajectories that violated the constraints for different discretizations. The reason why the method failed is most likely that the chosen grids introduced an error big enough in part of the space, that the optimal control inputs could not be identified anymore. This can be backed up by the fact that as the discretization becomes more dense the number of uncontrollable trajectories diminish. Also important that again at the last step the number of uncontrollable trajectories increase which could be explained by the fact that the system does not want to incur large control inputs, leading to some uncontrollable states. Compared to the 1D case, where we had fewer uncontrollable trajectories in the one step ahead predictions, it is important to note, that as the systems become more complex, a higher discretization is needed for the proposed method to be able to perform well. Both the forward minimization method and our benchmark were able to control all trajectories for all discretizations.

In Figure 4-10, we present the absolute error between the two methods and our benchmark, for the one step ahead control input, averaged over the initial conditions. We assume at each step the starting state to be the optimal one, so that we compare the two methods fairly. The highlighted regions represent the space in which $80 \%$ of the data points for each step lie.

Similarly to the 1D example the proposed method performs worse than the current one and the difference becomes smaller as we increase the discretization. Even in the denser grids, the proposed method performs two times worse than the current method. This can be attributed to the larger error of extracting the conjugate pairs compared to the cojugate function that was discussed in a previous section.


Figure 4-9: Number of invalid trajectories at each step in the horizon.

As is the case in the 1D system, we will also present the methods' performance over the entire time horizon. Again the proposed method failed to control all of the trajectories without violating the constraints. In Figure 4-11, we present the accumulated number of trajectories that broke the constraint, at each step, for different discretizations. For the smallest discretization of 11 points, $90 \%$ of the total initial conditions could not be controlled by the algorithm. As the discretization increases, the number of invalid trajectories declines steeply. This result clearly emphasizes the need for increasing the discretization of the grid, or finding a better way of choosing the dual grid during the conjugation. For the remaining plots, we will be including only the valid trajectories for the proposed method. Both the forward minimization method and the benchmark were able to control all initial conditions without violating the constraint.

In Figure 4-12, we present the accumulating cost of the three methods at each step averaged over the initial conditions. Additionally, in Figure 4-13, we present the absolute error of the corresponding trajectories resulting from applying the control inputs we calculated from each method. The results of these plots, are mostly unchanged from what we observed in the 1D example. Two things need to be mentioned. The error of the trajectory seems to be increasing for both methods early in the horizon. For more complex systems, we can observe that the effect of sub-optimal control inputs creates a larger snowballing effect in the states as the trajectories drift further from the optimal one earlier than the 1D case. The second observation is the performance of the proposed method in the smallest discretization. In Figure 4-13a, even when including only the trajectories that didn't violate the constraints, the method barely manages to control them. The error for the method is hovering around 0.25 which is enormous given that the state space spans the interval $[-1,1]$. The much bigger errors, result in high control inputs in an effort to control the system which results to the almost 5 times bigger average total cost of the trajectories over the other methods. The sparse discretization coupled with a sub-optimal choice for the dual grid makes the proposed method essentially impractical.


Figure 4-10: Absolute error of one step ahead optimal control inputs against the benchmark averaged over 100 initial conditions.


Figure 4-11: Number of invalid trajectories at each step in the horizon for the multistep implementation.


Figure 4-12: Average aggregate cost at each step for the control inputs calculated by the 3 different methods.


Figure 4-13: Trajectory error for each step averaged over 100 initial positions.

## Chapter 5

## Conclusions and Future Research

In this final chapter of the thesis, we will shortly discuss the work conducted and shortcomings of the proposed methods as well as try to provide any willing reader with a few research directions in order to improve the overall CDP algorithm.

## 5-1 Conclusions

In this thesis, several improvements of the algorithms proposed in [1] were developed. Starting, we considered the application of Kuhn's triangulation as an interpolation technique. The goal here was twofold: First, check whether we can preserve convexity in a convex extensible set of discrete points. We proved that it is not possible to do that every time. Second, we tested the method against the current interpolation method used by CDP, LERP, in order to see whether one method outperforms the other. A more accurate representation of the cost-to-go function will lead to smaller errors in the optimal control input and thus smaller costs for the trajectory of the system. In our testing, the two methods performed similarly with one outperforming the other in about $15 \%$ of the total cases. The difference between the two methods became smaller as the grid we used became more dense. Overall, the use of a different interpolation technique will provide the algorithm with more versatility as the user can choose between the two and use the best performing one for their specific use case.

Lastly, we considered the extraction of the optimal control input from within the CDP algorithm. To achieve that, we employed the conjugation's optimal pairs in order to create the optimal control law $u=u^{*}\left(y^{*}(x)\right)$, where our optimal control input $u$ is a function of the dual variable $y$, which in turn is a function of our current state $x$. We showed that extraction of the optimal control input is indeed feasible when dealing with continuous spaces, but depends highly on the discretization of the dual grid, when the spaces we are working with are discrete. Since we need to choose the dual grid before actually performing the conjugation, the actual optimal pairs may not be within that grid and thus we introduce an error creating a suboptimal set of pairs. We showed in our example that for the same discretization of the dual grid, the error introduced by the sub-optimal pairs is larger than the error of the conjugate
function itself. This is extremely important since the conjugate functions computed in [1] are currently used in order to extract the optimal control sequence. So for sub-optimal grids the current method employed by the CDP should outperform the proposed method. Another important observation is that for the same grid size, the error of the dual pairs is much larger outside the border of the optimal dual grid than inside it. This can be explained since every point outside the border of the optimal dual grid gets assigned to the same point in the primal grid, thus allowing the error to increase the further away we go. By having a larger dual grid than required the points outside the optimal dual grid are not useful and thus we make the effective region of the space more sparse, resulting in worse results when extracting the optimizer mappings. In our simulations, the same behaviour was again evident. The proposed method performed slighty worse than the current method for the one step ahead predictions of the control input. For the entire horizon the cost of the proposed method was much worse due to the accumulating error of the multistep implementation. By increasing the density of the grid, we observed that the both methods performed better and also closer to each other. Since both methods can theoretically extract the optimal control sequence, by increasing the grid closer to its continuous limit, the error between the two methods should shrink. To sum up, as it currently stands, the proposed method should not be used in the CDP algorithm, since the error it introduces is larger than the current method's, especially in smaller grids. Improving the selection of the dual grids during the conjugation steps of the algorithm could lead to significant improvements which could make the proposed method more advantageous, should that happen.

## 5-2 Future Research

Several improvements could be carried out in the works presented in the current thesis. These suggestions are pertaining to the work conducted during this thesis. For further research ideas one could read the relevant section of the original CDP paper.

Identification of system matrices. One possible research idea to extend the capabilities of the original algorithms is to introduce an identification technique that is able to generate a model. Since the original algorithms are model based, by employing an identification technique we can solve a larger number of problems, not necessarily model based. Especially important is the identification when dealing with noisy data, as the performance of the algorithm depends on the accuracy of the model.

Preserving Convexity. As neither the current method, LERP, nor the proposed one, Kuhn triangulation, can preserve convexity when extending a convex extensible set of points, further research is required in finding a computationally efficient interpolation method that could guarantee that.

Choosing the dual grid. As mentioned above the choice of the dual grid greatly impacts the error of the proposed method for extracting the optimal control input. By choosing more fitting grids the proposed method could result in significantly lower errors, which could make the method more attractive than the currently used one since its computational load is effectively non existent.

## CDP Algorithms

In this appendix we present the CDP algorithms that were developed in [1].

## Setting 1

```
Algorithm 1 Implementation of the d-CDP operator for Setting 1 [1]
Input: dynamics \(F_{s}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}, F_{i}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n x m} ;\) cost-to-go (at t+1) \(J: \mathbb{X}_{g} w R\); conjugate
    of stage cost \(C_{x}^{*}: \mathbb{X}_{g} x \mathbb{R}^{m} \rightarrow R\)
Output: cost-to-go (at t) \(\hat{\mathcal{T}}[J](x): \mathbb{X}_{g} \rightarrow R\)
    1: construct the grid \(\mathbb{Y}_{g}\);
    2: use LLT to compute \(J^{d *}: \mathbb{Y}_{g} \rightarrow R\);
    3: for each \(x \in \mathbb{X}_{g}\) do
        \(\phi_{x}(y) \leftarrow C_{x}^{*}\left(-F_{i}(x)+J^{d *}(y)\right)\) for \(y \in \mathbb{Y}_{g}\)
        \(\hat{\mathcal{T}}[J](x) \leftarrow \max _{y \in \mathbb{Y}_{g}}\left\{\left\langle F_{s}(x), y\right\rangle-p h i_{x}(y)\right\}\)
    end for
```


## Setting 2

```
Algorithm 2 Implementation of the d-CDP operator for Setting 2 [1]
Input: dynamics \(F_{s}: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}, B: \in \mathbb{R}^{n x m}\); cost-to-go (at t+1) \(J: \mathbb{X}_{g} \rightarrow R\); stage cost
    (state) \(C_{s}(x): \mathbb{X}_{g} x \mathbb{R}^{n} \rightarrow R\); conjugate of stage cost (input) \(C_{i}^{*}: \mathbb{R}^{m} \rightarrow R\); grid \(\mathbb{Z}_{g} \subset R\)
Output: cost-to-go (at t) \(\hat{\mathcal{T}}_{d}^{m}[J](x): \mathbb{X}_{g} \rightarrow R\)
    construct the grid \(\mathbb{Y}_{g}\);
    use LLT to compute \(J^{d *}: \mathbb{Y}_{g} \rightarrow R\) from \(J: \mathbb{X}_{g} \rightarrow \mathbb{R}\);
    \(\phi(y) \leftarrow C_{i}^{*}\left(-B^{T} y\right)+J J^{d *}(y)\) for \(y \in \mathbb{Y}_{g}\)
    use LLT to compute \(\phi^{d *}: \mathbb{Z}_{g} \rightarrow \mathbb{R}\) from \(\phi: \mathbb{Y}_{g} \rightarrow \mathbb{R}\)
    for each \(x \in \mathbb{X}_{g}\) do
        use LERP to compute \(\overline{\phi^{d * d}}\left(F_{s}(x)\right)\) from \(\phi^{d *}: \mathbb{Z}_{g} \rightarrow \mathbb{R}\)
        \(\hat{\mathcal{T}}_{d}^{m}[J](x) \leftarrow C_{s}(x)+\overline{\phi^{d * d}}\left(F_{s}(x)\right)\)
    end for
```


## Bibliography

[1] M. A. S. Kolarijani and P. M. Esfahani, "Fast approximate dynamic programming for input-affine dynamics," 2021.
[2] R. E. Bellman, Dynamic Programming. USA: Dover Publications, Inc., 2003.
[3] M. B. Haugh and L. Kogan, "Chapter 22 duality theory and approximate dynamic programming for pricing american options and portfolio optimization," in Financial Engineering (J. R. Birge and V. Linetsky, eds.), vol. 15 of Handbooks in Operations Research and Management Science, pp. 925-948, Elsevier, 2007.
[4] M. Uddin, M. Romlie, M. Abdullah, S. Abd Halim, A. Bakar, and T. Kwang, "A review on peak load shaving strategies," Renewable and Sustainable Energy Reviews, vol. 82, pp. 3323-3332, 112017.
[5] A. Oudalov, R. Cherkaoui, and A. Beguin, "Sizing and optimal operation of battery energy storage system for peak shaving application," in 2007 IEEE Lausanne Power Tech, pp. 621-625, 2007.
[6] Q. Wei, G. Shi, R. Song, and Y. Liu, "Adaptive dynamic programming-based optimal control scheme for energy storage systems with solar renewable energy," IEEE Transactions on Industrial Electronics, vol. 64, no. 7, pp. 5468-5478, 2017.
[7] R. Martins, H. C. Hesse, J. Jungbauer, T. Vorbuchner, and P. Musilek, "Optimal component sizing for peak shaving in battery energy storage system for industrial applications," Energies, vol. 11, no. 8, 2018.
[8] F. Borrelli, M. Baoti, A. Bemporad, and M. Morari, "Dynamic programming for constrained optimal control of discrete-time linear hybrid systems," Automatica, vol. 41, no. 10, pp. 1709-1721, 2005.
[9] Q. Wei, D. Liu, and H. Lin, "Value iteration adaptive dynamic programming for optimal control of discrete-time nonlinear systems," IEEE Transactions on Cybernetics, vol. 46, no. 3, pp. 840-853, 2016.
[10] A. Kleywegt, V. Nori, and M. Savelsbergh, "Dynamic programming approximations for a stochastic inventory routing problem," Transportation Science, vol. 38, pp. 42-70, 02 2004.
[11] D. Bertsekas and J. Tsitsiklis, Neuro-Dynamic Programming, vol. 27. 011996.
[12] D. Bertsekas, Dynamic programming and optimal control: Volume I, vol. 1. Athena scientific, 2012.
[13] R. S. Sutton and A. G. Barto, Reinforcement learning: An introduction. MIT press, 2018.
[14] C. J. Watkins and P. Dayan, "Q-learning," Machine learning, vol. 8, no. 3-4, pp. 279-292, 1992.
[15] J. N. Tsitsiklis and B. Van Roy, "An analysis of temporal-difference learning with function approximation," IEEE transactions on automatic control, vol. 42, no. 5, pp. 674-690, 1997.
[16] B. V. R. D.P. de Farias, "The linear programming approach to approximate dynamic programming," Operations Research, vol. 51, no. 6, pp. 850-865, 2003.
[17] F. Bach, "Max-plus matching pursuit for deterministic markov decision processes," 2019.
[18] E. Berthier and F. Bach, "Max-plus linear approximations for deterministic continuousstate markov decision processes," IEEE Control Systems Letters, vol. 4, no. 3, pp. 767772, 2020.
[19] R. T. ROCKAFELLAR, Convex Analysis. Princeton University Press, 1970.
[20] Y. Lucet, "Faster than the Fast Legendre Transform, the Linear-time Legendre Transform," Numerical Algorithms, vol. 16, pp. 171-185, Mar. 1997.
[21] S. Davies, "Multidimensional triangulation and interpolation for reinforcement learning," in Advances in Neural Information Processing Systems (M. C. Mozer, M. Jordan, and T. Petsche, eds.), vol. 9, MIT Press, 1997.
[22] D. W. Moore, Simplicial Mesh Generation with Applications. PhD thesis, USA, 1992. UMI Order No. GAX93-00795.
[23] K. Murota, "Discrete convex analysis," Mathematical Programming, vol. 83, no. 1, pp. 313-371, 1998.
[24] K. Murota, Recent Developments in Discrete Convex Analysis, pp. 219-260. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009.
[25] K. Murota, Discrete Convex Analysis: Monographs on Discrete Mathematics and Applications 10. USA: Society for Industrial and Applied Mathematics, 2003.


[^0]:    ${ }^{1}$ This column depicts the percentage of cases one method outperforms the other for the same discretization.

