Linear Algebra and its Applications 471 (2015) 693–729

Contents lists available at ScienceDirect

Linear Algebra and its Applications

www.elsevier.com/locate/laa

A max-plus based fundamental solution for a class of discrete time linear regulator problems ☆

Huan Zhang *, Peter M. Dower

Department of Electrical and Electronic Engineering, University of Melbourne, Melbourne, Victoria 3010, Australia

A R T I C L E   I N F O

Article history:
Received 7 January 2014
Accepted 12 January 2015
Available online 3 February 2015
Submitted by V. Mehrmann

MSC:
93C05
15A80
49L20
49J15

Keywords:
Max-plus fundamental solution
Linear quadratic regulator
Difference Riccati equations

A B S T R A C T

Efficient Riccati equation based techniques for the approximate solution of discrete time linear regulator problems are restricted in their application to problems with quadratic terminal payoffs. Where non-quadratic terminal payoffs are required, these techniques fail due to the attendant non-quadratic value functions involved. In order to compute these non-quadratic value functions, it is often necessary to appeal directly to dynamic programming in the form of grid- or element-based iterations for the value function. These iterations suffer from poor scalability with respect to problem dimension and time horizon. In this paper, a new max-plus based method is developed for the approximate solution of discrete time linear regulator problems with non-quadratic payoffs. This new method is underpinned by the development of new fundamental solutions to such linear regulator problems, via max-plus duality. In comparison with a typical grid-based approach, a substantial reduction in computational effort is observed in applying this new max-plus method. A number of simple examples are presented that illustrate this and other observations.

☆ This research is supported by grants FA2386-12-1-4084 from AFOSR and DP120101549 from Australian Research Council respectively.

* Corresponding author.
E-mail addresses: hzhang5@unimelb.edu.au (H. Zhang), pdower@unimelb.edu.au (P.M. Dower).

http://dx.doi.org/10.1016/j.laa.2015.01.008
0024-3795/

1. Introduction

After more than 40 years of study, the "linear quadratic regulator problem" (or LQR problem) remains ubiquitous in the field of optimal control [2,6]. Given a specific linear time invariant system, quadratic running payoff, and terminal payoff, the objective of the LQR (optimal control) problem is to determine a control sequence that (when applied to the linear system in question) maximizes the aggregated running and quadratic terminal payoffs over a specific (possibly infinite) time horizon. It is well known that the value function defined by the LQR problem is quadratic. The Hessian of this quadratic value function is either the solution of a difference (or differential) Riccati equation (DRE) in the finite horizon case, or the stabilizing solution of an algebraic Riccati equation (ARE) in the infinite horizon case. Solutions to either equation, and hence the corresponding LQR problem, can be computed very accurately and efficiently using existing numerical tools (for example, MATLAB™).

Both the DRE and ARE encode invariance of the space of quadratic functions (defined on the state space) with respect to the dynamic programming evolution operator associated with a quadratic running payoff and linear dynamics. Consequently, both equations are restricted in their application to problems with quadratic terminal payoffs. Where the terminal payoff employed is non-quadratic, the DRE/ARE solution path for the corresponding linear regulator problem is inherently invalid (as the corresponding value function involved need not be quadratic). Instead, it is necessary to appeal directly to the dynamic programming principle to obtain an iteration for the value function. This iteration is in general infinite dimensional, regardless of the state dimension. Consequently, approximate value function iterations employing state-space grids, basis functions, etc., arise out of necessity, but remain intrinsically limited in their application due to the curse-of-dimensionality. Consequently, where the time horizon is long or the state dimension high, the approximate solution of a linear regulator problem in the company of a non-quadratic terminal payoff remains a computationally expensive (and sometimes even prohibitive) exercise.

In this paper, a new computational method is developed for approximating the value function associated with a class of discrete time linear regulator problems in which the terminal payoff is non-quadratic. Motivated by recent related work [17,8,10,9], this new method relies on the development of a max-plus based fundamental solution for the class of linear regulator problems of interest. Using max-plus duality arguments [1,3,7,13,15–17], this fundamental solution captures the behaviour of the associated dynamic programming evolution operator, and is independent of the terminal payoff employed. By applying this fundamental solution to the terminal payoff associated with a specific linear regulator problem, the attendant value function (and hence the solution of this linear regulator problem) may be computed. Furthermore, by appealing to the algebraic structure of the fundamental solution, a substantial improvement in computation time relative to grid-based iterative methods can be achieved. This improvement is demonstrated via a number of computational examples. In addition, the limiting behaviour