Systems & Control Letters 106 (2017) 58–67

Contents lists available at ScienceDirect

Systems & Control Letters

journal homepage: www.elsevier.com/locate/sysconle

A general approach to the eigenstructure assignment for reachability and stabilizability subspaces

Lorenzo Ntogramatzidis *, Fabrizio Padula

Department of Mathematics and Statistics, Curtin University, Perth WA 6845, Australia

ARTICLE INFO

Article history:
Received 7 March 2017
Received in revised form 19 May 2017
Accepted 8 June 2017

Keywords:
LTI systems
Eigenstructure assignment

ABSTRACT

This paper is concerned with the problem of determining basis matrices for the supremal output-nulling, reachability and stabilizability subspaces, and the simultaneous computation of the associated friends that place the assignable closed-loop eigenvalues at desired locations. Our aim is to show that the Moore–Laub algorithm in Moore and Laub (1978) for the computation of these subspaces was stated under unnecessary restrictive assumptions. We prove the same result under virtually no system-theoretic hypotheses. This provides a theoretical foundation to a range of recent geometric techniques that are more efficient and robust, and as general as the standard ones based on the computation of sequences of subspaces.

1. Introduction

One of the central notions of geometric control theory is the one of output-nulling subspaces. In the past forty years, these subspaces and their duals have been found to play a fundamental role in an impressive number of control and estimation problems, including disturbance decoupling [1], non-interacting control [2,3], model matching [4], optimal control and filtering [5,6], unknown-input observation [7,8], fault detection [9], and patterned systems [10].

The three most important output-nulling subspaces used in all these problems (along with their duals) are:

- $\mathscr{V}^\star$, which represents the initial states of an LTI system for which there exists a control function that maintains the output identically at zero;
- $\mathscr{R}^\star$, which is the set of states that are reachable from the origin by means of a suitable control function that simultaneously maintains the output at zero;
- $\mathscr{V}_g^\star$, which represents the initial states for which there exists a control function that maintains the output identically at zero and at the same time ensures that the state of the system converges asymptotically to the origin.

A key feature of these subspaces is that the control functions that achieve the corresponding defining properties can always be expressed in terms of a static state feedback $u = F\,x$, where the feedback matrix F is referred to as a *friend* of the associated output-nulling subspace. Obtaining a basis for the subspace $\mathscr{V}_g^\star$ requires eigenspace computations [2,8]. On the contrary, basis matrices for the other two subspaces can be obtained in finite terms by resorting to sequences of subspaces that are guaranteed to converge in a finite number of steps, which is not greater than the order of the system. Once a basis has been obtained, the second problem is usually the calculation of a friend F that has desirable closed-loop properties. In fact, the choice of the friend F is not unique, and enables some parts of the closed-loop spectrum to be arbitrarily assigned. A fundamental result that showed that some alternative methods can be used to compute basis matrices for the aforementioned output-nulling subspaces is the one that will be referred to here as the *Moore and Laub algorithm* of [11]. This algorithm, unlike the classical methods, does not use a sequence of subspaces, but hinges on the calculation of null-spaces of the associated Rosenbrock system matrix pencil. It was shown in [11] that the use of this algorithm is the key to find significant computational improvements in the calculation of basis matrices for the main controlled invariant subspaces of geometric control.

In a recent paper [12], the algorithm of Moore and Laub was shown to be adaptable to the case in which the friend is sought to also assign the closed-loop eigenstructure that is *external* to $\mathscr{R}^\star$ (or $\mathscr{V}^\star$, or $\mathscr{V}_g^\star$), by deriving a parametric form for all the friends that assign any desired inner and outer closed-loop free spectrum. The consequent degrees of freedom can therefore be exploited to derive an algorithm that allows to obtain a friend that – in addition to placing the desired assignable closed-loop eigenvalues at arbitrary locations – also minimizes the Frobenius condition number of the matrix of closed-loop eigenvectors, which is a commonly used robustness measure, or that minimizes the Frobenius

* Corresponding author.
E-mail addresses: L.Ntogramatzidis@curtin.edu.au (L. Ntogramatzidis), Fabrizio.Padula@curtin.edu.au (F. Padula).

http://dx.doi.org/10.1016/j.sysconle.2017.06.003
0167-6911/

norm of the friend itself. It was also shown in [12] via extended Monte Carlo simulations that the method based on the parametric form derived from the Moore and Laub algorithm dramatically improves eigenvalue insensitivity to parameter uncertainties, with significantly smaller gain and vastly improved accuracy (by some orders of magnitude) than the friends obtained from the only other two publicly available MATLAB® toolboxes *GA* and *Linear Systems Toolkit* of [3] and [13], respectively.

The results in [11] have also been used in a number of contexts, including fault detection and tracking control, see [9,14].

Despite the wealth of results and contributions that in the past thirty years have cited [11] for its breakthrough in terms of the computation of $\mathscr{R}^\star$ and $\mathscr{V}^\star$, so far there have not been further developments on the method proposed in [11]. This is somewhat surprising, considering the fact that the main result of [11], which is Proposition 4, has been presented under a set of unnecessary restrictive assumptions. First, only the strictly proper case was addressed in [11]. Moreover, matrix B was assumed to be of full column rank, and the cardinality of the (distinct) set of eigenvalues $\mathscr{L}$ used to construct a basis for $\mathscr{R}^\star$ (which eventually appear as the closed-loop spectrum restricted to $\mathscr{R}^\star$) was assumed in [11] to be greater than or equal to the dimension of $\mathscr{R}^\star$. Furthermore, this set was also constrained not to contain elements whose real part is an invariant zero of the system. We will present and prove this result without the need for any of these assumptions. We will also address another delicate issue, which remained unsolved in [11], but which is crucial in the construction of an algorithm based on this result: the statement of Proposition 4 requires the knowledge of the dimension of $\mathscr{R}^\star$, which is the subspace whose basis we seek to construct. This dependence of the construction method on the size of $\mathscr{R}^\star$ creates a logical loop that raises some doubts on the effectiveness of Proposition 4 as an algorithm for the calculation of a basis for $\mathscr{R}^\star$. We show in this paper that, while the dimension of $\mathscr{R}^\star$ is used in the statement of this result, because such dimension provides an indication on the number of eigenvalues in the set of closed-loop eigenvalues to be assigned, this is not an obstacle to the construction of an algorithm that computes a basis for $\mathscr{R}^\star$ in finite terms.

In the last part of the paper, we remove the last restriction of [11] which requires $\mathscr{L}$ to contain distinct elements. We show in particular that the statement of [11, Proposition 4] still holds if the multiplicity of the values of $\mathscr{L}$ are such that the closed-loop matrix is non-defective. However, in the defective case, the procedure in [11, Proposition 4] cannot generate the entire $\mathscr{R}^\star$. We show here that there is no trivial generalization for the procedure of [11, Proposition 4] to address the defective case, in particular in the presence of Jordan chains related to the invariant zeros of the system, and we propose an algorithm to overcome this limitation. The importance of this procedure is not only theoretical. Indeed, a further advantage of employing the null-spaces of the Rosenbrock matrix for the calculation of output-nulling subspaces lies in the simultaneous calculation, as a by-product, of a corresponding friend that assigns the eigenstructure given in the (multi)set $\mathscr{L}$. Thus, a generalization of [11, Proposition 4] to the defective case yields an algorithm for the computation of a friend of $\mathscr{R}^\star$ (and, considering also the contribution of the invariant zeros, of $\mathscr{V}^\star$ and $\mathscr{V}_g^\star$) which assigns the eigenvalues of the closed-loop matrix restricted to $\mathscr{R}^\star$ with *any* desired multiplicity and any admissible Jordan structure. However, a second even more delicate issue will be addressed. While the construction of a spanning set for $\mathscr{R}^\star$ and $\mathscr{V}^\star$ is essentially the same in the case of distinct closed-loop eigenvalues, when we have repeated eigenvalues a fundamental difference arises in the way bases for $\mathscr{R}^\star$ and $\mathscr{V}^\star$ are obtained; this is due to the fact that the Jordan chains relative to the invariant zeros cannot be constructed starting from any vector of the corresponding eigenspace, but, in general, the chains need to start from only specific directions of that eigenspace. A method is presented to select the vector that can be used as starting point of a Jordan chain. This issue has not been considered in [15], where no distinction was made in the Jordan chains constructed from invariant zeros and from values different from invariant zeros. Therefore, the matrices H_i in [15, Theorem 4] are not completely free in general, but have to satisfy a further constraint (the one imposing that the vectors can indeed generate the entire Jordan chain) which will be discussed here.

The results presented in this paper also validate a number of geometric methods that in the past thirty years have been published, which hinge on the algorithm of [11] without acknowledging its limitations. In particular, this extension guarantees that the results in [15] indeed allow to compute a friend of $\mathscr{R}^\star$ and $\mathscr{V}^\star$ with the desired admissible Jordan structure. This paper provides the theoretical foundation for the numerical methods developed in [15]. These methods can be effectively used to compute output-nulling bases and friends also in the presence of non-trivial Jordan forms, despite the well-known numerical issues associated with the defective case.

Indeed, the determination of the Jordan form of a defective matrix is a well-known computational challenge. Nevertheless, the methods based on the calculation of the null-space of the reachability matrix pencil for the pole placement and of the Rosenbrock matrix pencil for the output nulling problem have been shown to produce numerically stable feedback matrices, [15,16].

While avoiding ill-conditioned problems may be possible for the assignable part of the closed-loop spectrum, situations involving defective zero structures require a general, systematic framework to deal with non-trivial Jordan forms.

In the proof of the exhaustiveness of the parameterization of the friends of $\mathscr{R}^\star$ of Theorem 1 in [12], the authors invoked [11, Proposition 4], but did not restrict their result to the same set of assumptions. The same issue arises in the proof of Theorem 3 of [15].

Notation. Throughout this paper, the symbol 0_q will stand for the origin of the vector space $\mathbb{R}^q$. For convenience, a linear mapping between finite-dimensional spaces and a matrix representation with respect to a particular basis are not distinguished notationally. The image and the kernel of matrix A are denoted by im A and ker A, respectively. The Moore–Penrose pseudo-inverse of A is denoted by $A^\dagger$. The spectrum of a square matrix A, denoted by $\sigma(A)$, is the multi-set of the eigenvalues of A counting multiplicities. Given a linear map $A : \mathscr{X} \longrightarrow \mathscr{Y}$ and a subspace $\mathscr{S}$ of $\mathscr{Y}$, the symbol $A^{-1}\,\mathscr{S}$ stands for the inverse image of $\mathscr{S}$ with respect to the linear map A, i.e., $A^{-1}\,\mathscr{S} = \{x \in \mathscr{X} \mid A\,x \in \mathscr{S}\}$. If $\mathscr{J} \subseteq \mathscr{X}$, the restriction of the map A to $\mathscr{J}$ is denoted by $A \mid \mathscr{J}$. If $\mathscr{X} = \mathscr{Y}$ and $\mathscr{J}$ is A-invariant, the eigenvalues of A restricted to $\mathscr{J}$ are denoted by $\sigma\,(A \mid \mathscr{J})$. If $\mathscr{J}_1$ and $\mathscr{J}_2$ are A-invariant subspaces and $\mathscr{J}_1 \subseteq \mathscr{J}_2$, the mapping induced by A on the quotient space $\mathscr{J}_2/\mathscr{J}_1$ is denoted by $A \mid \mathscr{J}_2/\mathscr{J}_1$, and its spectrum is denoted by $\sigma\,(A \mid \mathscr{J}_2/\mathscr{J}_1)$. The symbol $\oplus$ stands for the direct sum of subspaces. The symbol $\uplus$ denotes multiset aggregation, i.e., union with any common elements repeated. Given a map $A : \mathscr{X} \longrightarrow \mathscr{X}$ and a subspace $\mathscr{B}$ of $\mathscr{X}$, we denote by $\langle A \mid \mathscr{B} \rangle$ the smallest A-invariant subspace of $\mathscr{X}$ containing $\mathscr{B}$. The symbol i stands for the imaginary unit, i.e., $\mathrm{i} = \sqrt{-1}$. The symbol α^* denotes the complex conjugate of $\alpha \in \mathbb{C}$. Finally, given a matrix M, we denote by M_i its ith row and by M^j its jth column, respectively.

2. Geometric preliminaries

In what follows, whether the underlying system evolves in continuous or discrete time is irrelevant. Accordingly, we denote by $\mathbb{T}$ the time index set of any signal, on the understanding that this represents either $\mathbb{R}^+$ in the continuous time or $\mathbb{N}$ in the discrete