IEEE Control Systems Letters paper presented at
2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Guaranteed Stability Margins for Decentralized Linear Quadratic Regulators

Mruganka Kashyap and Laurent Lessard

***Abstract*— It is well-known that linear quadratic regulators (LQR) enjoy guaranteed stability margins, whereas linear quadratic Gaussian regulators (LQG) do not. In this letter, we consider systems and compensators defined over directed acyclic graphs. In particular, there are multiple decision-makers, each with access to a different part of the global state. In this setting, the optimal LQR compensator is dynamic, similar to classical LQG. We show that when sub-controller input costs are decoupled (but there is possible coupling between sub-controller state costs), the decentralized LQR compensator enjoys similar guaranteed stability margins to classical LQR. However, these guarantees disappear when cost coupling is introduced.**

I. Introduction

Multi-agent systems with communication constraints occur naturally in engineering applications, including bilateral teleoperation systems in remote robotic surgery and unmanned aerial vehicles (UAVs). For example, a swarm of UAVs could be deployed to survey an uncharted region or to optimize geographic coverage while combating forest fires. Information transfer within the swarm could be limited due to geographic constraints such as mountains blocking line-of-sight communications between certain UAVs.

It is known that certain decentralized information-sharing architectures lead to tractable optimal control problems [1], [2]. One such problem is *decentralized LQR* where the communication constraints have a poset-causal architecture [3], [4]. Although this is a state-feedback problem, the optimal decentralized controller is *dynamic* and has an observer-regulator structure reminiscent of output-feedback LQG regulators.

Robustness is an important aspect of controller design, because it ensures that the controller can effectively and reliably control a system in the presence of disturbances, plant uncertainty, or unmodeled dynamics. In the centralized case, LQR controllers enjoy guaranteed gain and phase margins [5], [6]. However, linear quadratic Gaussian (output feedback) regulators, have no robustness guarantees [7].

The robustness properties of decentralized LQR are not immediately apparent, since decentralized LQR shares commonalities with both centralized LQR (uses state feedback), and centralized LQG (optimal controller is dynamic). To the best of our knowledge, this is an open problem.

In this letter, we show that decentralized LQR enjoys similar stability margins to classical LQR if the input matrix (B) and control weighting matrix (R) are block-diagonal. We also show via counterexample that these assumptions are necessary.

In Sections II and III we review classical stability margins for LQR and more recent work on decentralized LQR synthesis. In Section IV we present our main results, and in Sections V and VI we present our counterexample and conclude.

II. Classical LQR stability margins

Consider the continuous-time linear time-invariant (LTI) dynamical system $\dot{x} = Ax + Bu$, where $x(t) \in \mathbb{R}^n$ and $u(t) \in \mathbb{R}^m$. The linear quadratic regulator (LQR) problem is to find the causal state-feedback policy that minimizes the quadratic cost

$$J = \int_0^\infty \big(x(t)^\mathsf{T} Q x(t) + u(t)^\mathsf{T} R u(t)\big)\,\mathrm{d}t. \tag{1}$$

Proposition 1: Suppose (A, B) is stabilizable, (Q, A) is detectable, and $Q \succeq 0$ and $R \succ 0$. The optimal LQR policy is $u(t) = Fx(t)$, where $F = -R^{-1}B^\mathsf{T}X$, and $X \succeq 0$ is the unique stabilizing solution to the algebraic Riccati equation $A^\mathsf{T}X + XA + Q - XBR^{-1}B^\mathsf{T}X = 0$.

We denote the optimal LQR gain from Theorem 1 using the notation $F := \mathrm{Ric}(A, B, Q, R)$. The optimal LQR controller is known to be inherently robust [8, §23] [9, §14.4]. In particular, if we define the loop gain $L(s) := F(sI-A)^{-1}B$, then the *Kalman inequality* holds:

$$(I - L(j\omega))^* R\,(I - L(j\omega)) \succeq R \quad \text{for all } \omega \in \mathbb{R}. \tag{2}$$

In the single-input case, $L(j\omega)$ is a scalar and the Kalman inequality reduces to $|1-L(j\omega)| \geq 1$. This can be interpreted as the open-loop Nyquist plot of $-L$ (negative feedback) lying outside the disk centered at $(-1, 0)$ with radius 1. This implies that the LQR compensator has gain margin $\frac{1}{2} < k < \infty$ and phase margin $-60° < \phi < 60°$.

Alternatively, a sufficient condition for robust stability can be expressed in terms of the perturbation itself [6].

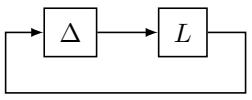

Fig. 1. Perturbed feedback interconnection. $L(s) := F(sI - A)^{-1}B$ is the loop gain for a standard LQR feedback controller $u(t) = Fx(t)$.

This material is based upon work supported by the National Science Foundation under Grant No. 2136317.

M. Kashyap is with the Department of Electrical and Computer Engineering at Northeastern University, Boston, MA 02115, USA. kashyap.mru@northeastern.edu

L. Lessard is with the Department of Mechanical and Industrial Engineering at Northeastern University, Boston, MA 02115, USA. l.lessard@northeastern.edu

Lemma 2: Consider the setting of Theorem 1, and let $L(s) := F(sI - A)^{-1}B$ be the LQR-optimal loop gain. The interconnected system of Fig. 1 is well-posed and internally stable for all LTI systems Δ that satisfy

$$\Delta(j\omega)^* R + R\Delta(j\omega) \succ R \quad \text{for all } \omega \in \mathbb{R}. \tag{3}$$

Proof: Invert (2) and apply the matrix inversion lemma, which yields $(I + H(j\omega))^* R (I + H(j\omega)) \preceq R$, where we defined the closed-loop map $H(s) := F(sI - A - BF)^{-1}B$. This is equivalent to $\|R^{1/2}(I + H)R^{-1/2}\|_\infty \leq 1$. Then, perform a loop-shifting transformation to Fig. 1 to obtain Fig. 2. Apply the small gain theorem [9, Thm. 9.1] to conclude that the interconnection is well-posed and internally stable for all LTI systems Δ satisfying $\|R^{1/2}(I - \Delta^{-1})R^{-1/2}\|_\infty < 1$, which is equivalent to (3). ■

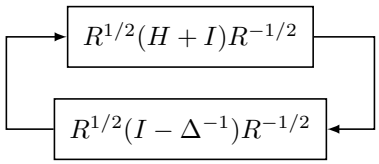

Fig. 2. Transformation of Fig. 1. $H(s) := F(sI - A - BF)^{-1}B$ is the closed-loop map for a standard LQR feedback controller.

Theorem 2 allows us to specialize the previous gain and phase margin results derived from the Kalman inequality to the case where each input channel is separately perturbed.

Corollary 3: Consider the setting of Theorem 2. Partition the input $u(t)$ into subvectors of dimension $m_1 + \cdots + m_N = m$. If we assume R and Δ are block-diagonal and partitioned conformally to the partition of $u(t)$, i.e.,

$$R = \begin{bmatrix} R_1 & & 0 \\ & \ddots & \\ 0 & & R_N \end{bmatrix} \quad \text{and} \quad \Delta = \begin{bmatrix} \Delta_1 & & 0 \\ & \ddots & \\ 0 & & \Delta_N \end{bmatrix},$$

then the interconnected system of Fig. 1 is well-posed and internally stable for all independent LTI perturbations of the blocks of $u(t)$ satisfying $\Delta_i(j\omega)^* R_i + R_i \Delta_i(j\omega) \succ R_i$ for $i = 1, \ldots, N$ and for all $\omega \in \mathbb{R}$. In particular, each input block independently has gain margin $\frac{1}{2} < k_i < \infty$ and phase margin $-60^\circ < \phi_i < 60^\circ$.

In Theorem 3, the assumption that R is block-diagonal is necessary. It is possible to construct systems where a non-diagonal R leads to closed loops that be destabilized by arbitrarily small perturbations in a single channel [6, Ex. 3.1].

Similar robustness results to Theorem 2 have been derived for discrete time [10] and for the case with cross-product cost terms [11], though these cases generally have weaker robustness guarantees. There are also negative results; when R is full, the independent perturbation result of Theorem 3 no longer holds [6, Ex. 3.1]. Finally, there are no guaranteed stability margins for LQG compensators [7].

III. Decentralized LQR control

We consider the problem setting studied in [3], [4], which is an LQR problem structured according to a directed acyclic graph (DAG). Specifically, we assume the setting in Theorem 1, but we partition the state as $x = \begin{bmatrix} x_1^\mathsf{T} & \cdots & x_N^\mathsf{T} \end{bmatrix}^\mathsf{T}$ and similarly for the input u. We also partition A and B as $N \times N$ block matrices conforming to the partitions of x and u.

There is an underlying DAG on the nodes $1, \ldots, N$, which are assumed to be ordered according to the partial ordering of the DAG. The matrices A and B have a block-sparsity pattern that conforms to the adjacency matrix of the transitive closure of the DAG. Consider for example the 4-node DAG in Fig. 3.

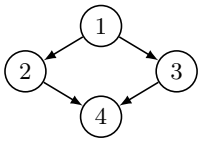

$$S = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}$$

Fig. 3. Example of a 4-node directed acyclic graph (DAG), the adjacency matrix of its transitive closure is S, shown on the right.

The associated dynamical system would have the structure

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \\ \dot{x}_4 \end{bmatrix} = \begin{bmatrix} A_{11} & 0 & 0 & 0 \\ A_{21} & A_{22} & 0 & 0 \\ A_{31} & 0 & A_{33} & 0 \\ A_{41} & A_{42} & A_{43} & A_{44} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} + \begin{bmatrix} B_{11} & 0 & 0 & 0 \\ B_{21} & B_{22} & 0 & 0 \\ B_{31} & 0 & B_{33} & 0 \\ B_{41} & B_{42} & B_{43} & B_{44} \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \end{bmatrix}$$

There are no assumptions on the cost matrices, so all states and inputs may be coupled through Q and R, respectively.

Definition 4: The ancestors of node i, denoted $\mathcal{A}(i)$, is the set of all nodes j for which there exists a directed path from j to i, including node i. Similarly, the descendants of node i, denoted $\mathcal{D}(i)$, is the set of all nodes j for which there exists a directed path from i to j, including i. We also use these sets as a matrix subscripts to indicate the submatrix formed by selecting the corresponding block rows and columns.

For the example of Fig. 3, we have $\mathcal{D}(2) = \{2, 4\}$ and $\mathcal{A}(3) = \{1, 3\}$, which defines the block submatrices

$$A_{\mathcal{D}(2)} = \begin{bmatrix} A_{22} & 0 \\ A_{24} & A_{44} \end{bmatrix} \quad \text{and} \quad B_{\mathcal{A}(3)} = \begin{bmatrix} B_{11} & 0 \\ B_{31} & B_{33} \end{bmatrix}.$$

What makes the problem *decentralized* is that each u_i only has access to the past history of the ancestors of node i. For the example of Fig. 3, this means the u_i take the form

$$u_1 = \mathcal{K}_1(x_1), \qquad u_2 = \mathcal{K}_2(x_1, x_2),$$
$$u_3 = \mathcal{K}_3(x_1, x_3), \qquad u_4 = \mathcal{K}_4(x_1, x_2, x_3, x_4),$$

where the $\mathcal{K}_i$ are causal maps. In general, decentralized problems with LQG assumptions need not have linear optimal controllers [12]. However, when the plant and controller are structured according to a DAG as above, the optimal controller is linear [2] and finding the optimal linear controller may be cast as a convex optimization problem [13].

Explicit closed-form solutions have been obtained for this decentralized LQR problem using a state-space approach [3], [14] and poset-based approach [4]. Similar explicit solutions exist for LQG (output-feedback) versions of this problem [14]–[17] and also with time delays [18], [19].

The optimal controller for the decentralized LQR problem described above has the following structure [3].

Proposition 5: Consider the decentralized LQR problem. Suppose (A_i, B_i) is stabilizable for $i = 1, \dots, N$ and (Q, A) is detectable. Let $F_i := \mathrm{Ric}(A_{\mathcal{D}(i)}, B_{\mathcal{D}(i)}, Q_{\mathcal{D}(i)}, R_{\mathcal{D}(i)})$. The optimal decentralized LQR controller has closed-loop dynamics and associated optimal policy given by

$$\begin{aligned} \dot{\xi}_i &= (A_{\mathcal{D}(i)} + B_{\mathcal{D}(i)} F_i)\xi_i \\ u_i &= \sum_{j \in \mathcal{A}(i)} I_{i,\mathcal{D}(j)} F_j \xi_j \qquad \text{for } i = 1, \dots, N \end{aligned}$$

where $I_{i,\mathcal{D}(j)}$ is the block-row of the identity matrix $I_{\mathcal{D}(j)}$ associated with node i.

If we include zero-mean process noise in the plant dynamics that is independent between the different nodes of the DAG, then $\xi_i = \mathbf{E}\big(x_{\mathcal{D}(i)} \mid x_{\mathcal{A}(i)}\big) - \mathbf{E}\big(x_{\mathcal{D}(i)} \mid x_{\mathcal{A}(i)\setminus\{i\}}\big)$, so ξ_i is an estimation correction in updating the estimate of the descendants once the current node i is included.

The optimal decentralized controller from Theorem 5 is linear, but unlike the classical centralized case in Theorem 1, it is also *dynamic*. The decentralized LQR controller bears a resemblance to the optimal LQG controller because its states are estimates of plant states. The main difference is that the strict descendants of node i are not observable, so rather than using an observer such as a Kalman filter, the state estimates are formed via *prediction* [20, §IV.D].

IV. Main Results

For the optimal decentralized LQR controller described in Theorem 5, there is no large Kalman inequality of the form (2). Instead, we have N separate Kalman inequalities

$$(I - L_i(j\omega))^* R_{\mathcal{D}(i)} (I - L_i(j\omega)) \succeq R_{\mathcal{D}(i)} \quad \forall \omega \in \mathbb{R}, \quad (4)$$

corresponding to the N separate centralized LQR subproblems that make up the optimal decentralized controller.

Consequently, there is no apparent way to leverage the small gain theorem as in the proof of Theorem 2. Instead, we show that if we assume B and R are block-diagonal, we can prove a result similar to Theorem 3 for block-diagonal perturbations.

Theorem 6: Consider the decentralized LQR problem and its optimal controller, described in Section III and Theorem 5, respectively, and let L_{dec} be the optimal loop gain.

Further suppose that R and B are block-diagonal with block sizes corresponding to the partitions of $x(t)$ and $u(t)$. The interconnected system of Fig. 4 is well-posed and internally stable for all independent LTI perturbations of the blocks of $u(t)$ satisfying the following for all $i = 1, \dots, N$.

$$\Delta_i(j\omega)^* R_i + R_i \Delta_i(j\omega) \succ R_i \quad \text{for all } \omega \in \mathbb{R}. \quad (5)$$

Remark 7: Theorem 6 looks similar to Theorem 3, but L_{dec} is now the more complicated loop gain for the optimal *decentralized* LQR controller. Unlike Theorem 3, Theorem 6 makes the additional assumptions that B and R are block diagonal. In Section V, we show that these assumptions are necessary, but we argue that they are not restrictive in many cases of practical interest.

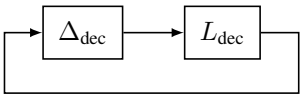

Fig. 4. Perturbed feedback interconnection. L_{dec} is the loop gain for the optimal decentralized LQR feedback controller described in Theorem 5 and $\Delta_{\text{dec}} = \mathrm{diag}\{\Delta_i\}$ is a block-diagonal LTI perturbation.

Proof: We take an approach similar to the proof of Theorem 2, except we use a more general version of the small gain theorem for structured uncertainty, and additional steps are required to combine the N separate Kalman inequalities into something we can use. Start by rewriting the closed-loop map of the optimal decentralized LQR controller from Theorem 5 as:

$$H_{\text{dec}} = 1_{\text{dec}}^{\mathsf{T}} \bar{F} (sI - \bar{A} - \bar{B}\bar{F})^{-1} \hat{B}$$

where we defined:

$$\begin{aligned} 1_{\text{dec}}^{\mathsf{T}} &:= \begin{bmatrix} I_{1,\mathcal{D}(1)} & \cdots & I_{N,\mathcal{D}(N)} \end{bmatrix} \\ \bar{A} &:= \mathrm{diag}\{A_{\mathcal{D}(i)}\} \\ \bar{F} &:= \mathrm{diag}\{F_i\} \\ \bar{B} &:= \mathrm{diag}\{B_{\mathcal{D}(i)}\} \\ \hat{B} &:= \begin{bmatrix} e_1 e_1^{\mathsf{T}} B_{\mathcal{D}(1)} \\ \vdots \\ e_N e_N^{\mathsf{T}} B_{\mathcal{D}(N)} \end{bmatrix}, \end{aligned}$$

where e_i is the i-th column of the identity matrix of size n. Since B is block-diagonal, we have $\hat{B} = \mathrm{diag}\{B_{\mathcal{D}(i)} e_i\}$. So we can rewrite the closed-loop map as $H_{\text{dec}} = 1_{\text{dec}}^{\mathsf{T}} \bar{H} E$, where we defined:

$$\bar{H} := \mathrm{diag}\{H_i\}, \quad E := \mathrm{diag}\{e_i\}, \quad \bar{R} := \mathrm{diag}\{R_{\mathcal{D}(i)}\},$$
$$H_i := F_i (sI - A_{\mathcal{D}(i)} - B_{\mathcal{D}(i)} F_i)^{-1} B_{\mathcal{D}(i)}.$$

Note that H_i is the closed-loop map for the separate LQR problem associated with $\mathcal{D}(i)$ defined in Theorem 5.

Now perform the same loop-shifting transformation as in the proof of Theorem 2 to Fig. 4 to obtain Fig. 5. Since R and Δ_{dec} are block-diagonal, the uncertainty block in Fig. 5 is also block-diagonal. Our goal is to apply the structured small gain theorem [9, Thm. 11.8], which is a generalization of the small gain theorem that applies when the uncertainty is structured.

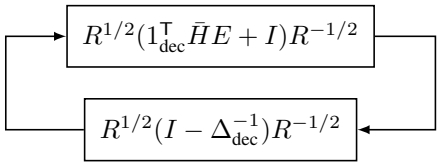

Fig. 5. Transformation of Fig. 4. $\bar{H} := \mathrm{diag}\{H_i\}$ is the block-diagonal concatenation of the closed-loop maps associated with the N centralized LQR sub-problems that make up the optimal decentralized LQR controller.

To this end, we state an intermediate lemma, which relates the structured singular value of the optimal closed-loop map to the N separate Kalman inequalities (4).

Lemma 8: Consider the setting of Theorem 6, where $\bar{H}$, H_i, E, and 1_{dec} are defined as above. The following inequality holds:

$$\sup_{\omega\in\mathbb{R}} \mu_\Delta\Big(R^{1/2}(1_{\text{dec}}^\mathsf{T}\bar{H}(j\omega)E+I)R^{-1/2}\Big) \leq \max_{1\leq i\leq N} \big\|R_i^{1/2}e_i^\mathsf{T}(I+H_i)e_iR_i^{-1/2}\big\|_\infty,$$

where $\mu_\Delta(\cdot)$ denotes the structured singular value corresponding to the block-diagonal structure of Δ_{dec}.

Proof: Let $M := R^{1/2}(1_{\text{dec}}^\mathsf{T}\bar{H}(j\omega)E+I)R^{-1/2}$. Since the plant and controller each have transfer functions structured according to the adjacency matrix S of the transitive closure of the associated DAG, they form an algebra. Consequently, all products, inverses, and linear fractional transformations preserve the structure, and in particular, so does the closed-loop map H_{dec}. Therefore, M has a block-sparsity structure conforming to S. Since the nodes are assumed to be ordered according to the partial ordering of the DAG, S is lower-triangular and so M is block-lower triangular.

Let $\mathbf{\Delta} := \{\text{diag}\{\Delta_i\} \mid \Delta_i \in \mathbb{C}^{m_i\times m_i}\}$. For any $\Delta\in\mathbf{\Delta}$,

$$\det(I-M\Delta) = \prod_{i=1}^N \det(I-M_{ii}\Delta_i)$$

and we can simplify M_{ii} based on the definition as

$$\begin{aligned} M_{ii} &= e_i^\mathsf{T}\left(R^{1/2}(1_{\text{dec}}^\mathsf{T}\bar{H}(j\omega)E+I)R^{-1/2}\right)e_i \\ &= R_i^{1/2}e_i^\mathsf{T}(1_{\text{dec}}^\mathsf{T}\bar{H}(j\omega)E+I)e_iR_i^{-1/2} \\ &= R_i^{1/2}e_i^\mathsf{T}(H_i(j\omega)+I)e_iR_i^{-1/2}. \end{aligned} \tag{6}$$

By the definition of the structured singular value,

$$\begin{aligned} \mu_\Delta(M) &= \frac{1}{\min\{\|\Delta\| \mid \det(I-M\Delta)=0, \Delta\in\mathbf{\Delta}\}} \\ &= \frac{1}{\min\{\|\Delta\| \mid \det(I-M_{ii}\Delta_i)=0 \text{ for some } i\}} \\ &\leq \frac{1}{\min_i \min\{\|\Delta_i\| \mid \det(I-M_{ii}\Delta_i)=0\}} \\ &= \max_{1\leq i\leq N}\|M_{ii}\|. \end{aligned}$$

The last step follows from the fact that $\mu_{\Delta_i}(M_{ii}) = \|M_{ii}\|$ because Δ_i is unstructured. Substituting in M_{ii} from (6) and taking the supremum over $\omega\in\mathbb{R}$ completes the proof. ∎

Inverting the Kalman inequalities in (4) and converting them into $\mathcal{H}_\infty$ norms as in the proof of Theorem 2, we obtain

$$\big\|R_{\mathcal{D}(i)}^{1/2}\big(I+H_i\big)R_{\mathcal{D}(i)}^{-1/2}\big\|_\infty \leq 1 \quad \text{for } i=1,\ldots,N.$$

Since for any matrix $M\in\mathbb{C}^{p\times q}$, the (spectral) norm of M is lower-bounded by the norm of any submatrix of M, have a similar inequality for $\mathcal{H}_\infty$ norms, and together with the fact that R is block-diagonal, we deduce that

$$\big\|R_i^{1/2}e_i^\mathsf{T}\big(I+H_i\big)e_iR_i^{-1/2}\big\|_\infty \leq \big\|R_{\mathcal{D}(i)}^{1/2}\big(I+H_i\big)R_{\mathcal{D}(i)}^{-1/2}\big\|_\infty$$

The two above inequalities together with Theorem 8 imply that

$$\sup_{\omega\in\mathbb{R}} \mu_\Delta\Big(R^{1/2}(1_{\text{dec}}^\mathsf{T}\bar{H}(j\omega)E+I)R^{-1/2}\Big) \leq 1.$$

We can now apply the structured small gain theorem [9, Thm. 11.8] and conclude that the interconnection of Fig. 5 is well-posed and stable whenever $\Delta_{\text{dec}} = \text{diag}\{\Delta_i\}$ satisfy

$$\big\|R^{1/2}(I-\Delta_{\text{dec}}^{-1})R^{-1/2}\big\|_\infty < 1.$$

Due to the block-diagonal structure of the uncertainty, this is equivalent to

$$\big\|R_i^{1/2}(I-\Delta_i^{-1})R_i^{-1/2}\big\|_\infty < 1 \quad \text{for } i=1,\ldots,N$$

which is equivalent to (5). ∎

Equipped with Theorem 6, we can specialize the decentralized LQR robustness result to the case where each input channel is perturbed using either a pure gain or a pure phase shift. This leads us to a decentralized version of Theorem 3.

Corollary 9: Consider the decentralized LQR setting of Theorem 6. Each input $u_i(t)$ independently has gain margin $\frac{1}{2} < k_i < \infty$ and phase margin $-60^\circ < \phi_i < 60^\circ$.

V. Discussion

Theorem 6 provides conditions for the robust stability of the optimal decentralized linear quadratic regulator, under the additional assumptions that B and R are block-diagonal and different perturbations are applied to each input u_i.

The assumption that B and R are block-diagonal is critical. We will demonstrate using a simple numerical example that the gain margin $\frac{1}{2} < k_i < \infty$ established in Theorem 9 no longer applies when either B or R is not block-diagonal.

Consider a two-node DAG with graph $1\to 2$ and global plant dynamics given by

$$\dot{x} = \begin{bmatrix}1 & 0\\ 1 & 1\end{bmatrix}x + \begin{bmatrix}1 & 0\\ \beta & 1\end{bmatrix}u, \tag{7}$$

cost matrices $Q = \left[\begin{smallmatrix}3 & 1\\ 1 & 3\end{smallmatrix}\right]$ and $R = \left[\begin{smallmatrix}100 & \rho\\ \rho & 100\end{smallmatrix}\right]$. We use the perturbation $\Delta_{\text{dec}} = \left[\begin{smallmatrix}k & 0\\ 0 & 1\end{smallmatrix}\right]$ with $k\in\mathbb{R}$, so node 1 is perturbed by a static scalar gain while node 2 remains unperturbed. The perturbed closed-loop matrix is given by

$$A_{\text{CL}} := \begin{bmatrix} 1+kF_1^{11} & kF_1^{12} & 0\\ 1+\beta F_1^{11}+F_1^{21} & 1+\beta F_1^{12}+F_1^{22} & 0\\ 1+k\beta F_1^{11}+F_1^{21} & F_1^{22}-F_2+k\beta F_1^{12} & 1+F_2\end{bmatrix}$$

where F_1^{ij} and F_2 are given by

$$\begin{bmatrix}F_1^{11} & F_1^{12}\\ F_1^{21} & F_1^{22}\end{bmatrix} = \text{Ric}\left(\begin{bmatrix}1 & 0\\ 1 & 1\end{bmatrix}, \begin{bmatrix}1 & 0\\ \beta & 1\end{bmatrix}, \begin{bmatrix}3 & 1\\ 1 & 3\end{bmatrix}, \begin{bmatrix}100 & \rho\\ \rho & 100\end{bmatrix}\right)$$

$$F_2 = \text{Ric}(1,1,3,100) \approx -2.0149.$$

The gain margin of input u_1 is the range of values of k for which A_{CL} is Hurwitz.

We ran two experiments. First, we assumed a diagonal R and triangular B, so we fixed $\rho = 0$ and varied β. Fig. 6 (top) shows a plot of the pairs (β, k) for which A_{CL} is Hurwitz (shaded in blue). When $\beta = 0$, we confirm the

result of Theorem 9; the system is stable for $\frac{1}{2} < k < \infty$, which corresponds to $k > -6\,\text{dB}$ on the plot. But when $\beta \neq 0$, violating the requirement that B be block-diagonal, we observe a severe deterioration in the gain margin.

For the second experiment, we assumed a full R but diagonal B, so we fixed $\beta = 0$ and varied ρ. Fig. 6 (bottom) shows a plot of the pairs (ρ, k) for which A_{CL} is Hurwitz (shaded in blue). As in the previous example, we confirm the result of Theorem 9 when $\rho = 0$, but we observe deterioration for some nonzero choices of ρ.

The matrices B and R are block-diagonal in many cases of practical interest. For example, consider multi-agent systems, such as drones flying in formation or a platoon of vehicles. In these cases, each control input affects a separate agent, so B is block-diagonal. Also, the total input cost is typically the sum of input costs for each agent, with no coupling. So R is block-diagonal as well.

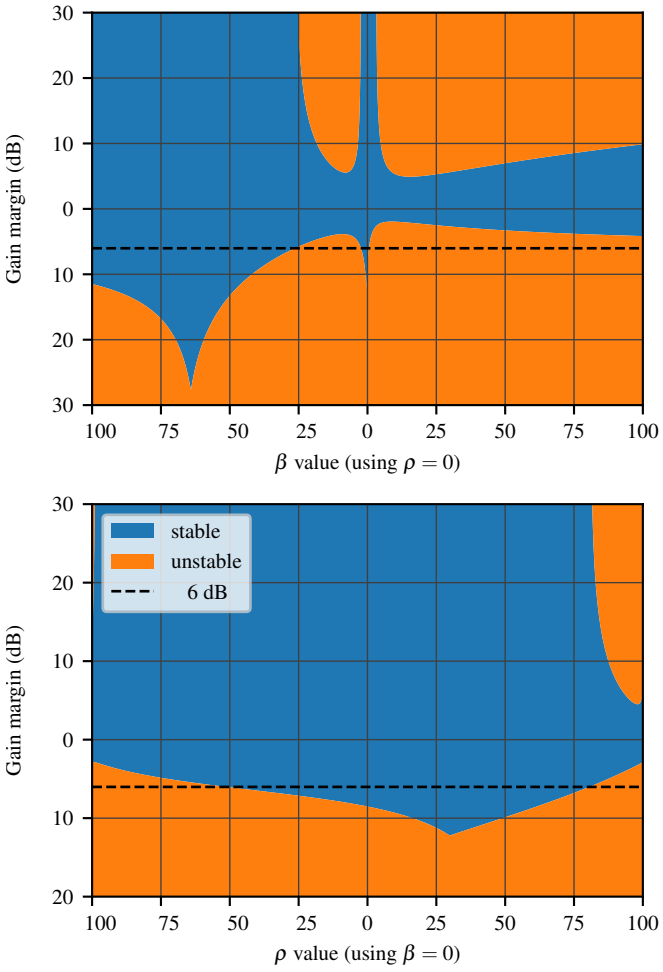

Fig. 6. Stability margins for the decentralized LQR example with the dynamics of Eq. (7). The input u_1 is perturbed by a factor of k. The top panel uses $\rho = 0$ (diagonal R) and the blue region shows the (β, k) that yield a stable closed loop, with k expressed in decibels (dB). The bottom panel uses $\beta = 0$ (diagonal B) and the blue region shows the (ρ, k) that yield a stable closed loop. When B and R are block-diagonal ($\rho = \beta = 0$), we recover Theorem 9, which ensures a gain margin $\frac{1}{2} < k < \infty$. In other words, $k > -6\,\text{dB}$.

VI. Conclusion

We studied the robustness of optimal decentralized LQR controllers when the plant and controller are structured according to a directed acyclic graph. Specifically, we established that when the B and R matrices are block-diagonal and different LTI perturbations are applied to each input, the controlled system enjoys the same stability margins as in the classical (centralized) LQR case. This is an interesting result because the optimal decentralized LQR controller is dynamic, much like an output-feedback LQG controller, yet LQG controllers have no stability margins.

While this letter only studied the case of LTI perturbations, our approach can be generalized to nonlinear input perturbations, analogous to the results obtained in [5].

References

[1] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *IEEE Trans. Automat. Contr.*, vol. 50, no. 12, pp. 1984–1996, 2005.

[2] Y.-C. Ho and K. Chu, "Information structure in dynamic multi-person control problems," *Automatica*, vol. 10, no. 4, pp. 341–351, 1974.

[3] J. Swigart and S. Lall, "Optimal controller synthesis for decentralized systems over graphs via spectral factorization," *IEEE Trans. Automat. Contr.*, vol. 59, no. 9, pp. 2311–2323, 2014.

[4] P. Shah and P. A. Parrilo, "$\mathcal{H}_2$-optimal decentralized control over posets: A state-space solution for state-feedback," *IEEE Trans. Automat. Contr.*, vol. 58, no. 12, pp. 3084–3096, 2013.

[5] M. Safonov and M. Athans, "Gain and phase margin for multiloop LQG regulators," *IEEE Trans. Automat. Contr.*, vol. 22, no. 2, pp. 173–179, 1977.

[6] N. Lehtomaki, N. Sandell, and M. Athans, "Robustness results in linear-quadratic Gaussian based multivariable control designs," *IEEE Trans. Automat. Contr.*, vol. 26, no. 1, pp. 75–93, 1981.

[7] J. C. Doyle, "Guaranteed margins for LQG regulators," *IEEE Trans. Automat. Contr.*, vol. 23, no. 4, pp. 756–757, 1978.

[8] J. P. Hespanha, *Linear systems theory, second edition*. Princeton university press, 2018.

[9] K. Zhou, J. C. Doyle, and K. Glover, *Robust and optimal control*. Prentice-Hall, Inc., 1996.

[10] U. Shaked, "Guaranteed stability margins for the discrete-time linear quadratic optimal regulator," *IEEE Trans. Automat. Contr.*, vol. 31, no. 2, pp. 162–165, 1986.

[11] D. Chung, T. Kang, and J. G. Lee, "Stability robustness of LQ optimal regulators for the performance index with cross-product terms," *IEEE Trans. Automat. Contr.*, vol. 39, no. 8, pp. 1698–1702, 1994.

[12] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM Jour. Contr.*, vol. 6, no. 1, pp. 131–147, 1968.

[13] M. Rotkowitz and S. Lall, "Decentralized control information structures preserved under feedback," in *IEEE Conf. Decis. Contr.*, vol. 1, 2002, pp. 569–575.

[14] J.-H. Kim and S. Lall, "Explicit solutions to separable problems in optimal cooperative control," *IEEE Trans. Automat. Contr.*, vol. 60, no. 5, pp. 1304–1319, 2015.

[15] L. Lessard and S. Lall, "Optimal control of two-player systems with output feedback," *IEEE Trans. Automat. Contr.*, vol. 60, no. 8, pp. 2129–2144, 2015.

[16] T. Tanaka and P. A. Parrilo, "Optimal output feedback architecture for triangular LQG problems," in *Amer. Contr. Conf.*, 2014, pp. 5730–5735.

[17] M. Kashyap and L. Lessard, "Explicit agent-level optimal cooperative controllers for dynamically decoupled systems with output feedback," in *IEEE Conf. Decis. Contr.*, 2019, pp. 8254–8259.

[18] ——, "Agent-level optimal LQG control of dynamically decoupled systems with processing delays," in *IEEE Conf. Decis. Contr.*, 2020, pp. 5980–5985.

[19] A. Lamperski and L. Lessard, "Optimal decentralized state-feedback control with sparsity and delays," *Automatica*, vol. 58, pp. 143–151, 2015.

[20] P. Shah and P. A. Parrilo, "An optimal controller architecture for poset-causal systems," in *IEEE Conf. Decis. Contr.*, 2011, pp. 5522–5528.