A Problem in Positive Systems Stability Arising in Topology Control*

Florian Knorn, Rade Stanojevic, Martin Corless and Robert Shorten

Abstract We present a problem in the stability of switched positive systems that arises in network topology control. Preliminary results are given that guarantee stability of a network topology control problem under certain assumptions. Roughly speaking, these assumptions reduce the underlying stability problem to a nonlinear consensus problem with a driving term, that eventually becomes a Lur'e problem. Simulation results are given to illustrate our algorithm. While these results indicate that our assumptions can be removed, a proof of the general stability problem remains open.

1 Introduction

Recent years have witnessed a growing interest in the control community in problems that arise when dynamic systems evolve over graphs. While the most high profile of these applications are in consensus applications such as formation flying and synchronisation problems, [4, 8, 11], many other applications have arisen where the manner in which network topologies change affect the performance of algorithms that are run over these networks. In such applications, an essential requirement is that the topology of the graph be such that some properties required to support communication and control are satisfied, the most basic of these being that the network is connected. Considerations of this kind have given rise to the emerging field of network topology control. Clearly, graph connectivity is an essential component in

Florian Knorn, Rade Stanojevic, Robert Shorten
Hamilton Institute, National University of Ireland Maynooth, Maynooth, Co. Kildare, Ireland,
E-mail: { florian.knorn , rade.stanojevic , robert.shorten } @nuim.ie

Martin Corless
School of Aeronautics & Astronautics, Purdue University, West Lafayette, IN 47907, USA,
E-mail: corless@purdue.edu

situations where a group of network nodes must work together, in a decentralised manner, to achieve some global task. This issue of graph connectivity is therefore very important and has achieved much attention in various contexts recently.

In this paper we describe a recently proposed decentralised topology control algorithm [6] and suggest a simple way of adding weights to the states. This algorithm was posed to overcome some common assumptions in topology control (namely, that the underlying graph is symmetric). It exploits the fact that the rate of convergence of certain algorithms evolving over a graph is a good proxy for graph connectivity. Furthermore, this rate of convergence can be estimated in a decentralised manner, and can therefore be used to regulate graph connectivity. Under the assumption that the estimation problem and the control problem operate on different time-scales, stability can be demonstrated using elementary arguments. In particular, we show that the feedback system reduces to a consensus problem with an input, and it eventually becomes a scalar nonlinear system that can be analysed in a Lur'e problem framework. Simulations are presented to illustrate the validity of our results. These results also indicate that the feedback system is stable even when a separation of time-scales is not present. This latter problem in positive systems remains open and is posed in the concluding remarks of the paper.

2 Main results

In the context of this paper and the topology control problem discussed in [6], we are interested in the following type of n-dimensional positive systems with an input term:

$$\mathbf{x}(k+1) = f\big(\mathbf{x}(k),k\big) + u\big(\mathbf{x}(k),k\big)\mathbf{1}, \qquad \mathbf{x}(0) = \mathbf{x}_0, \qquad k = 0,1,2,\dots$$

where, $\mathbf{x}(k) \in \mathbb{R}^n_+$ are the states, $f : \mathbb{R}^n_+ \times \mathbb{Z} \to \mathbb{R}^n_+$ is a continuous vector-valued function, and $u(\mathbf{x}(k),k) \in \mathbb{R}$ is an input term, that, through $\mathbf{1} = (1 \ \dots \ 1)^\mathsf{T}$, is thus equally added to all states and that we assume to be such that the system's states do not leave the positive orthant. We would then like to investigate under which conditions the system's states approach each other over time and, in the limit, eventually all take the same value — which may be time-varying, depending on the input term. More formally, we are looking for conditions such that $\lim_{k\to\infty}\big(x_i(k)-x_j(k)\big) = 0$ for all $i,j \in \{1,\dots,n\}$.

2.1 Affine case

Before stating our more general result we would like to present a more easily established result which can be obtained when the function f takes a particular linear form: $f\big(\mathbf{x},k\big) = \mathbf{P}(k)\mathbf{x}$ where $\mathbf{P}(k) \in \mathbb{R}^{n\times n}$ is a sequence of primitive, row-stochastic

(and thus non-negative) matrices with strictly positive main diagonal entries. This special form is often encountered in distributed averaging or consensus applications, for instance.

Theorem 1. *Let* $\mathbf{P}(k) \in \mathbb{R}^{n \times n}$ *be a sequence of matrices taken from a finite set of primitive, row-stochastic matrices with strictly positive main diagonal entries, and* $u(\mathbf{x}(k), k)$ *a sequence of real, non-negative numbers. If* $\mathbf{x}(k) \in \mathbb{R}_+^n$ *evolves for some* $\mathbf{x}(0) = \mathbf{x}_0 \in \mathbb{R}_+^n$ *according to*

$$\mathbf{x}(k+1) = \mathbf{P}(k)\mathbf{x}(k) + u(\mathbf{x}(k), k)\mathbf{1} \tag{1}$$

where $\mathbf{1} = (1 \ldots 1)^\mathsf{T}$, *the elements of* $\mathbf{x}(k)$ *will approach each other over time, that is* $\lim_{k \to \infty} x_i(k) - x_j(k) = 0$ *for all* $i, j \in \{1, \ldots, n\}$.

Proof. Given in [6]. □

2.2 General case

The previous result can be extended to classes of nonlinear consensus operators using the recent results of [7]. Borrowing its notation, we get:

Theorem 2. *Let* $\mathscr{G}(k) = (\mathscr{V}, \mathscr{A}(k))$ *be a sequence of strongly connected graphs*[2], $u(\mathbf{x}(k), k)$ *a sequence of finite real numbers and* f *a map on* $\mathbb{R}^n \times \mathbb{Z}$ *satisfying the following. Associated to each directed graph* $\mathscr{G} = (\mathscr{V}, \mathscr{A})$ *with node set* $\mathscr{V} = \{1, \ldots, n\}$, *each node* $i \in \mathscr{V}$ *and each state* $\mathbf{x} \in \mathbb{R}_+^n$, *there is a compact set* $\mathscr{E}_i(\mathscr{A})(\mathbf{x}) \subset \mathbb{R}$ *satisfying:*

1. $f_i(\mathbf{x}, k) \in \mathscr{E}_i(\mathscr{A}(k))(\mathbf{x})$ $\forall k \in \mathbb{N}$ $\forall \mathbf{x} \in \mathbb{R}_+^n$,
2. $\mathscr{E}_i(\mathscr{A})(\mathbf{x}) = \{x_i\}$ *whenever the states of node* i *and its neighbouring nodes* j *are all equal,*
3. $\mathscr{E}_i(\mathscr{A})(\mathbf{x})$ *is contained in the relative interior of the convex hull of the states of node* i *and its neighbouring nodes* j *whenever the states of node* i *and its neighbouring nodes* j *are not all equal,*
4. $\mathscr{E}_i(\mathscr{A})(\mathbf{x})$ *depends continuously on* $\mathbf{x}$, *that is, the set-valued function* $\mathscr{E}_i(\mathscr{A}) : \mathbb{R}_+^n \rightrightarrows \mathbb{R}_+$ *is continuous.*

Then, if the states $\mathbf{x}(k) \in \mathbb{R}_+^n$ *evolve for some initial condition* $\mathbf{x}(0) = \mathbf{x}_0 \in \mathbb{R}_+^n$ *according to*

$$\mathbf{x}(k+1) = f(\mathbf{x}(k), k) + u(\mathbf{x}(k), k)\mathbf{1}$$

where $\mathbf{1} = (1 \ldots 1)^\mathsf{T}$, *the elements of* $\mathbf{x}(k)$ *will approach each other over time, i. e.* $\lim_{k \to \infty} (x_i(k) - x_j(k)) = 0$ *for all* $i, j \in \{1, \ldots, n\}$.

Proof. Given in [6]. □

---

[2] That is, there is a directed path connecting any two nodes in the network.

Remark. Put simply, the theorem's four conditions require that the updated state of each node must be a strict convex combination of its own and its neighbours' states, and that the update function must be continuous.

3 Application of main result

3.1 Distributed averaging and topology control

Consensus or distributed averaging algorithms have been the subject of an inordinate amount of attention in the past decade, as they arise in applications such as distributed sensing, clock synchronisation, flocking, or fusion of Kalman filter data; see for instance [2, 4, 9, 10]. Since the rate at which these algorithms converge strongly depends on structural properties of the network of nodes they are run on, it is an interesting problem to try to somehow regulate the topology of the graph in order to ultimately control the speed at which consensus will be achieved. But as control usually requires some form of measurement and feedback of the quantity of interest, in this case one would need to be able to determine the level of connectivity.

While the primary focus of the present paper is neither on the properties nor the dynamics of consensus algorithms, we recall that the second eigenvalue in magnitude of the averaging matrix[3] determines the rate at which the nodes in a network achieve consensus. Classically, the second smallest eigenvalue of the *Laplacian* (or *transition Laplacian*) matrix of a graph has been used as an algebraic measure for connectivity, [1, 3]. However, Laplacians are usually only defined for symmetric graphs, a restriction that we would like to avoid. In that regard, the second eigenvalue of an averaging matrix is also an excellent candidate measurement to indicate the degree of connectivity of an entire graph (whether the underlying graph is directed or not). It also has the added benefit of being able to be estimated locally.

In [6] we describe several methods of estimating this important, global quantity in a distributed way. With the algorithms provided therein, each node in a network is able to estimate the second eigenvalue using only local, readily available information. In wireless networks (or, on a more abstract level, geometric graphs), this would offer the abovementioned possibility to control or maintain a certain level of interconnectedness: Each node could reduce or expand its communication radius if the connectivity is estimated to be larger or smaller than required (as decreasing or increasing this radius will lead to reducing or increasing the number of neighbours, hence changing connectivity). That such a strategy is well posed is evident and follows from the basic observation that if all nodes increase their communication radii sufficiently, then the graph will eventually achieved the desired level of connectedness. Let us investigate this control application more concretely in the following.

---

[3] Many distributed averaging algorithms can be written as $\mathbf{x}(k+1) = \mathbf{P}\mathbf{x}(k)$ where $\mathbf{x}(k)$ is the vector containing the states of all the nodes in the network, at time k. The row-stochastic, so-called *averaging matrix* $\mathbf{P}$ describes how each node averages its own value with that of its neighbors.

3.2 Control strategy

Given a wireless network, we wish to adjust the communication radius of each node in the network $r_1, \dots, r_n > 0$ using the estimates of the second eigenvalue in magnitude of the averaging matrix, λ, with the ultimate objective of regulating λ to some neighbourhood of a target value; namely so that $|\lambda - \lambda^*| < \varepsilon$ for some $\lambda^* \in (0, 1)$ and $\varepsilon > 0$. Since there will always be more than one set of communication radii $\{r_1, \dots, r_n\}$ that will guarantee this objective, we shall propose a control law that guarantees that the closed loop algorithm converges to a *unique*, single radius used by all nodes. Although this additional requirement is made to facilitate analytical tractability (that is, uniqueness of the solution), it can also be motivated from a practical standpoint: Having all nodes use the same broadcast radius helps to achieve similar battery lifetimes of the nodes, which is desirable in many applications. However, our framework is sufficiently general to allow other quantities of interest to be included in the control law design (but the convergence proofs will change accordingly). For instance, relaxing the requirements on the communication radii, one may require all nodes to have an equal number of neighbours (which would also yield a unique radius distribution).

To achieve this, we propose updating the individual node radii using a convex combination of their neighbours' radii, plus an input term that depends on the estimated second largest eigenvalue, that is we feed back of the current level of connectivity.[4] Specifically, we suggest the following decentralised control law

$$\mathbf{r}(k+1) = \mathbf{P}_c(k)\mathbf{r}(k) + \eta \left[\lambda(k) - \lambda^*\right]\mathbf{1} \tag{2}$$

for some initial, strictly positive radius distribution $\mathbf{r}(0) = \mathbf{r}_0$ which guarantees a strongly connected graph. Here $\mathbf{P}_c(k)$ is a sequence of primitive, row-stochastic averaging matrices on the graphs induced by $\mathbf{r}(k)$, $\lambda(k)$ is the magnitude of the second largest eigenvalue of the averaging matrix $\mathbf{P}$ as in Footnote 3 for the graph topology at time k, and $\eta > 0$ is a suitable control gain. To be fully precise, we could write $\lambda(\mathbf{r}(k))$ to highlight that the second eigenvalue is ultimately a function of the topology of the graph, which in turn is dependent on $\mathbf{r}(k)$. Unfortunately both dependencies are rather complex in nature and hard to determine or express analytically. However, it can be shown that $\lambda(k)$ may be treated as a sector bounded nonlinearity so that it can be treated in a Lur'e framework, [5]: Both Theorems 1 and 2 will guarantee that the control law forces all the radii, over time, to a common value. In other words, (2) will eventually become a *scalar* relation, so that the stability and convergence properties of the controlled system will eventually be governed by the scalar, positive system

$$x(k+1) = x(k) + u(x(k), k)$$

---

[4] Note that we assume a certain separation of time scales between the estimation and the control scheme, i. e. we assume that the estimators have successfully converged to an exact estimate that is hence common to all nodes in the network.

Since the properties of such systems are well understood, the above theorems offer interesting possibilities for the design of further control laws. This also allows us to determine how the control gain η must be chosen so that the closed loop system is stable, which is reported in depth in [6]. Note again, that any other consensus scheme (to which Theorem 2 can be applied) may be used as well. Also, we would like to stress that the proposed controller is decentralised in that each node only requires the radius information of its neighbours — information that can easily be broadcast along the communication that is necessary to run the algorithm used for estimating $\lambda(k)$ in the first place.

Remark. Let us comment on possible connectivity issues when using the above control law. When λ^* is chosen very close to one, it may be possible that in some iteration the control law would adjust (reduce) the communication radii so much that network becomes disconnected. This can either be prevented by using a much smaller control gain than necessary for stability (which guarantees that λ^* is approached without overshoot), or by introducing a "minimum radius" that the nodes' radii are not allowed to fall below and that is large enough to guarantee that the graph always remains strongly connected.

3.3 Weighting

In the scheme presented above, all nodes eventually reach a common radius. Now, imagine a setting where some nodes (say, nodes 17 and 25) are equipped with a longer-lasting power supply than others. In that case, these "special" nodes should be allowed to use a larger broadcast radius relative to the consensus value of the other, "ordinary" nodes. This would be an example of a situation were a certain "weighting" is applied to the states of each node. We shall now see that this can easily be incorporated in our set-up, without changing the convergence proofs.

Let $\mathbf{W} := \text{diag}\{w_i\}$ with $w_i > 0$ for $i = 1, \ldots, n$ be the $n \times n$ diagonal matrix with positive entries w_i along its main diagonal, and let $\tilde{\mathbf{r}} := \mathbf{W}^{-1}\mathbf{r}$. Then, run the control strategy for the "auxiliary" states $\tilde{\mathbf{r}}$ — which will converge to a common value — but recover and utilise the "weighted" radii using $\mathbf{r} = \mathbf{W}\tilde{\mathbf{r}}$. In the example above, this would mean that setting $w_{17} = w_{25} = 2$, and $w_i = 1$ for all other nodes.

Remark. The proposed weighting could also be used in a slightly more elaborate way. For instance, a node's weight could be made a function of the remaining battery power, such that it is decreased over time as its battery is getting emptied. This way, nodes with little remaining battery power are allowed to use smaller radii than others so that they can "survive" a little bit longer: As their weight decreases, so will their communications radius, relative to the other nodes in the network. However, one will need to assure that the radius is not decreased too much so that the node disconnects from the network (see the remark at the end of the previous subsection).

4 Simulation Results

Let us now present some simulations that demonstrate our results. They all show experiments on networks with 200 nodes and randomly distributed initial radii in the $[0.05, 0.95]$ interval,[5] in which the second largest eigenvalue in magnitude was regulated to some desired value. Depicted are the evolution over time of the second eigenvalue together with the nodes' radii.
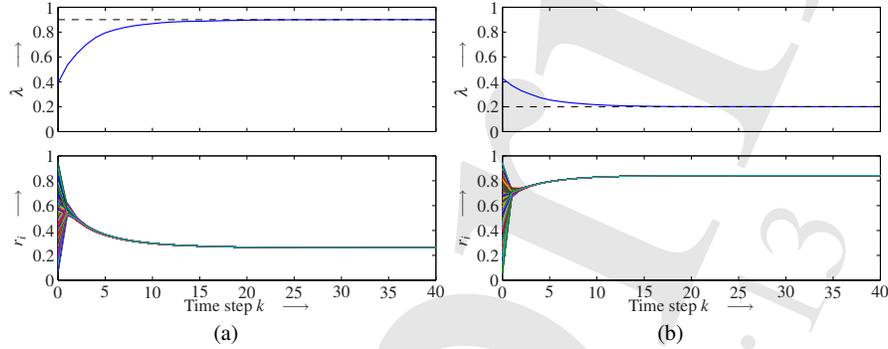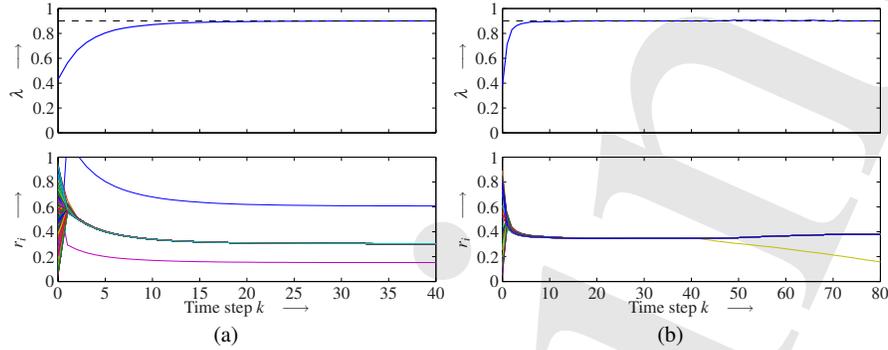

Fig. 1 Evolution of the second eigenvalue $\lambda(k)$ of the averaging matrix (upper subplots) and the individual nodes' communication radii $r_i(k)$ (lower plots) in two networks of 200 nodes for (a) $\lambda^* = 0.9$ and (b) $\lambda^* = 0.2$.

Both Figures 1(a) resp. (b) show situations where the nodes were to achieve $\lambda^* = 0.9$ resp. $\lambda^* = 0.2$ using a common communication radius. It can be seen that at first the radii converge to a common value and then, on a slower time scale, change such that the second eigenvalue reaches the desired value.

Figure 2 shows two examples where the nodes' states were weighted. In (a), we simulate the setting where two nodes are equipped with different power supplies than the others. This presents an application of weighting the states as mentioned in Subsection 3.3. We picked two nodes which we wanted to use twice resp. half the radius as the other nodes in the network. This was achieved by setting the corresponding weights to 2 resp. 0.5. As can be seen in the plot, the second largest eigenvalue of the network converges quickly to its desired value of $\lambda^* = 0.9$, and the nodes' radii all converge to a common value but for the two special nodes of different weighting.

An example for the remark at the end of Subsection 3.3 is given in Figure 2(b). Whilst again regulating the second largest eigenvalue in magnitude, we started to successively reduce one node's weight starting at time $k = 40$. The plots show that the desired level of connectivity is, again, quickly achieved and maintained throughout. However, after $k = 40$ one node's radius decreases bit by bit whereas the other

---

[5] We deliberately chose different initial radii to show that consensus is achieved on these values.

Fig. 2 Evolution of $\lambda(k)$ and the individual nodes' radii $r_i(k)$ in two networks of 200 resp. 50 nodes as the second eigenvalue is regulated to $\lambda^* = 0.9$. In (a) two nodes where to have twice resp. half the radius as their peers; in (b) one node's weight was successively reduced starting at $k = 40$.

nodes' radii all commonly increase, slightly, to counter the effect of the reduction in radius of the special node.

5 Conclusion and future directions

In the context of consensus algorithms, we presented two theorems that provide conditions for the convergence of the states to a common value, even when there are inputs to the system. We also suggested a simple modification that allows for different weightings to be applied to the states. We then used these results to control the topology of wireless networks or geometric graphs in general. The proposed decentralised control law, which adjusts the communication radii of the nodes so that the overall network achieves a predefined level of connectivity, poses such a consensus problem with inputs, and possible weighting of the states.

This leads us to the more general, open problem of finding consensus conditions for systems of the type

$$\mathbf{x}(k+1) = f\big(\mathbf{x}(k), k\big) + g\big(\mathbf{x}(k), \tilde{\boldsymbol{\lambda}}(k), k\big)$$

where f is some convex function of the system's states, and g is a function of the (local) states and the (local) estimates $\tilde{\lambda}_i(k)$ of the second largest eigenvalue of the averaging matrix of the graph. These systems are encountered when we drop the assumption of separation of time scales of estimation and control scheme, or in certain distributed optimisation problems, when g represents the derivative of some convex cost function.

References

1. Chung, F.R.K.: Spectral Graph Theory. No. 92 in CBMS Regional Conference Series in Mathematics. American Mathematical Society, Province, RI, USA (1997)
2. Estrin, D., Girod, L.D., Pottie, G.J., Srivastava, M.: Instrumenting the world with wireless sensor networks. In: Proc. of the Int. Conf. on Acoustics, Speech, and Signal Processing, vol. 4, pp. 2033–2036. Salt Lake City, UT, USA (2001)
3. Fiedler, M.: Algebraic connectivity of graphs. Czech. Math. J. **23**(98), 298–305 (1973)
4. Jadbabaie, A., Lin, J., Morse, A.S.: Coordination of groups of mobile autonomous agents using nearest neighbor rules. IEEE Trans. Autom. Control **48**(6), 988–1001 (2003)
5. Khalil, H.K.: Nonlinear Systems. Macmillan Publishing Co., New York, NY, USA (1992)
6. Knorn, F., Stanojevic, R., Corless, M., Shorten, R.: A framework for decentralised feedback connectivity control with application to sensor networks. To appear in Int. J. Control (2009)
7. Moreau, L.: Stability of multiagent systems with time-dependent communication links. IEEE Trans. Autom. Control **50**(2), 169–182 (2005)
8. Olfati-Saber, R., Murray, R.M.: Consensus problems in networks of agents with switching topology and time-delays. IEEE Trans. Autom. Control **49**(9), 1520–1533 (2004)
9. Reynolds, C.W.: Flocks, herds, and schools: A distributed behavioral model. In: Proc. of the 14th Annual Conf. on Computer Graphics and Interactive Techniques, pp. 25–34. Anaheim, CA, USA (1987)
10. Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., Shochet, O.: Novel type of phase transition in a system of self-driven particles. Phys. Rev. Lett. **75**(6), 1226–1229 (1995)
11. Zavlanos, M.M., Pappas, G.J.: Controlling connectivity of dynamic graphs. In: Proc. of the Joint 44th IEEE Conf. on Decision and Control, and the European Control Conf., pp. 6388–6393. Seville, Spain (2005)