Convexity of Optimal Control over Networks with Delays and Arbitrary Topology

Michael Rotkowitz[1] Randy Cogill[2] Sanjay Lall[3]

International Journal of Systems, Control, and Communications, Vol. 2, No. 1-3, pp. 30–54, 2010.

Abstract

We consider the design of optimal controllers for a networked system (or a spatio-temporal system), where the dynamics of each subsystem may affect those of other subsystems with some propagation delays, and the controllers may communicate with each other with some transmission delays. We show that if a simple condition holds, then the optimal control problem may be recast as a convex optimisation problem. This is shown to unify and broadly generalise the class of such systems amenable to convex synthesis. When we consider the special case of spatially invariant systems, this again broadly generalises the known class of tractable problems.

Keywords: decentralised control; distributed control; networked control; convex optimisation; quadratic invariance

1 Introduction

We consider the problem of multiple subsystems, each with its own controller, such that the dynamics of each subsystem may effect those of other subsystems with some propagation delay, and the controllers may communicate with each other with some transmission delays. We seek to synthesize linear controllers to minimize a closed-loop norm for the entire interconnected system. This is an optimal decentralized control problem which is difficult in general, and there is no known tractable solution for arbitrary propagation and transmission delays. This paper states simple conditions on the delays such that this optimal control problem may be cast as a convex optimization problem.

It has been shown for general decentralized control that a property called quadratic invariance allows the optimal control problem to be recast as a convex optimization problem [7]. We thus achieve our characterization of delays which allow for convex synthesis by testing for quadratic invariance.

We find that if the transmission delays satisfy the triangle inequality, and if the propagation delay between any pair of subsystems is at least as large as the transmission delay between those subsystems, then the problem is quadratically invariant. In other words,

[1]Department of Electrical and Electronic Engineering,
The University of Melbourne, Parkville VIC 3010, Australia,
Email: mcrotk@unimelb.edu.au
[2]Department of Systems and Information Engineering,
The University of Virginia, Charlottesville, VA 22904, U.S.A.,
Email: rcogill@virginia.edu
[3]Department of Aeronautics and Astronautics,
Stanford University, Stanford CA 94305-4035, U.S.A,
Email: lall@stanford.edu

if data can be transmitted faster than dynamics propagate along any link, then optimal controller synthesis may be formulated as a convex optimization problem.

It is important to note the extreme generality of this framework and of this result. It holds for discrete-time systems and continuous-time systems. It holds for any norm that we wish to minimize. It does not assume that the dynamics of any subsystem are the same as those of any other, and they may all be completely different types of objects. Most importantly, the delay between any two subsystems is not assumed to have any relationship whatsoever to other delays in the system. They may be assigned independently for each link. Only in the examples do we assume otherwise.

We then consider spatio-temporal systems, with sensor measurements and control actions at each point in a spatial domain. We find that when viewed from the proper general framework, where delays and subsequent effects between any two points may be arbitrarily assigned, then the approach and the results for networks of subsystems with delays can be seamlessly extended. Thus we find that if the transmission delays satisfy the triangle inequality, and if the propagation delay between any pair of points is at least as large as the transmission delay between those points, then the problem is quadratically invariant.

Given this general result, we then view spatially invariant systems as a special case, and what falls out almost immediately is a broad generalisation of the class of spatially invariant systems for which one can convexify the optimal distributed control problem.

Some of these results, in particular, the crucial role of the triangle inequality, were first presented in [5]. It has since been shown that the same conditions also allow for the parameterization of all stabilizing causal (possibly nonlinear, possibly time-varying) controllers over a network [4].

1.1 Prior Work

A vast amount of prior work on optimal control over networks assumes that the actions of any subsystem have no effect on the dynamics of other subsystems. For a few other specific structures, tractable methods have been found. One of the first problems of this nature to be studied was the one-step delayed information sharing problem. This problem assumes that each subsystem has a controller that can see its own output immediately, and can see outputs from all other subsystems after a delay of one time step. This problem has long been known to admit tractable solutions [9], and has also been studied more recently in an LFT framework [8]. An interesting class of spatio-temporal systems which allow for convex synthesis of optimal controllers was identified in [1], and named funnel causal systems. One of the tractable structures discussed in [3] involved evenly spaced subsystems which can pass measurements on at the same speed that the dynamics propagate, and [7] included a similar class of evenly spaced systems where the bound was found such that if the communication speed exceeded that bound the problem was amenable to convex synthesis.

These results are all unified and generalized by the simple conditions found in this paper.

1.2 Outline

In Section 2, we state some preliminaries and notation, define the propagation and transmission delays, explain why we may assume that the transmission delays satisfy the triangle inequality, formulate the problem we wish to solve, and give an overview of results on quadratic invariance, in particular, that it allows convex synthesis of optimal linear decentralized controllers.

In Section 3, we consider the problem of control over networks. This includes the main result that being able to communicate faster than the dynamics propagate between any pair of nodes results in the optimal decentralised control problem being convex, provided that the transmission delays satisfy the triangle inequality.

In Section 4 we consider spatio-temporal systems. We first derive an analogous result for this class of problems, and then consider the special case of spatially invariant systems, resulting in a generalisation of that class which can be convexified.

We make some concluding remarks in Section 5.

2 Preliminaries

We make use of the standard L_p Banach spaces equipped with the usual p-norm, and the extended space

$$L_{2e} = \{f : \mathbb{R}_+ \to \mathbb{R} \mid f_T \in L_2 \text{ for all } T \in \mathbb{R}_+\}$$

We use similar notation for discrete time. As is standard, we extend the discrete-time Banach spaces ℓ_p to the extended space

$$\ell_e = \{f : \mathbb{Z}_+ \to \mathbb{R} \mid f_T \in \ell_\infty \text{ for all } T \in \mathbb{Z}_+\}$$

Note that in discrete time, all extended spaces contain the same elements, since the common requirement is that the sequence is finite at any finite index. This motivates the abbreviated notation of ℓ_e.

We also make use of $\mathcal{L}_e$ when we want to overload our notation to simultaneously state something for L_{2e} and ℓ_e.

Delay. We define Delay$(\cdot)$ for a causal operator as the smallest amount of time in which an input can affect its output. For any causal $H : \mathcal{L}_e^m \to \mathcal{L}_e^m$,

$$\begin{aligned}\text{Delay}(H) = \inf\{\tau \geq 0 \mid z_1(T+\tau) \neq z_2(T+\tau), z_1 = H(w_1), z_2 = H(w_2),\\ w_1, w_2 \in \mathcal{L}_e^m, w_1(t) = w_2(t) \text{ for all } t \leq T\}\end{aligned}$$

and if $H = 0$, we consider its delay to be infinite.

When H is time-invariant, we may choose $T = 0$, and when H is linear, we may choose $w_1 = 0$, so for a causal Linear Time-Invariant (LTI) $H : \mathcal{L}_e^m \to \mathcal{L}_e^n$,

$$\text{Delay}(H) = \inf\{\tau \geq 0 \mid z(\tau) \neq 0, z = H(w), w \in \mathcal{L}_e^m, w(t) = 0 \text{ for all } t \leq 0\}$$

Given an impulse response h which characterizes the map H, we can then also give the delay as

$$\text{Delay}(H) = \operatorname{ess\,inf}\{\tau \geq 0 \mid h(\tau) \neq 0\}$$

Note that we then have the following inequalities for the delays of a composition or an addition of operators:

$$\begin{aligned}\text{Delay}(AB) &\geq \text{Delay}(A) + \text{Delay}(B)\\ \text{Delay}(A+B) &\geq \min\{\text{Delay}(A), \text{Delay}(B)\}\end{aligned}$$

LFT. We suppose that we have a generalized plant $P : \mathcal{W} \times \mathcal{U} \to \mathcal{Z} \times \mathcal{Y}$ partitioned as

$$P = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & G \end{bmatrix}$$

We define the ***closed-loop map*** by

$$f(P, K) = P_{11} + P_{12}K(I - GK)^{-1}P_{21}$$

The map $f(P, K)$ is also called the (lower) ***linear fractional transformation*** (LFT) of P and K. Note that we abbreviate $G = P_{22}$, since we will refer to that block frequently, and so that we can refer to its subdivisions without ambiguity. This interconnection is shown in Figure 1.

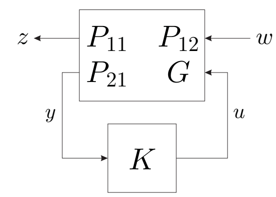

Figure 1: Linear fractional interconnection of P and K

Subsystems. We suppose that there are n subsystems, each with its own controller, which can be represented as nodes on a graph, such as in Figure 2. Associated with each subsystem $i \in 1, \ldots n$ is then a set of possible control actions $\mathcal{U}_i = L_{2e}^{q_i}$ (or $\ell_e^{q_i}$), and a set of possible output measurements $\mathcal{Y}_i = L_{2e}^{m_i}$ (or $\ell_e^{m_i}$). All of the results of this paper hold for continuous time or discrete time, and we will hereafter mostly refrain from stating both. We also define the overall control space $\mathcal{U} = \mathcal{U}_1 \times \cdots \mathcal{U}_n$ and the overall measurement space $\mathcal{Y} = \mathcal{Y}_1 \times \cdots \mathcal{Y}_n$.

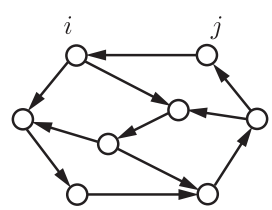

Figure 2: Arbitrary pair of nodes, arbitrary network

For any pair of nodes $i, j \in 1, \ldots n$, we then have a mapping from the control action at node j to the measurement at node i given as $G_{ij} : \mathcal{U}_j \to \mathcal{Y}_i$, and a part of controller i to be designed using the available information from node j, given as $K_{ij} : \mathcal{Y}_j \to \mathcal{U}_i$.

We can amalgamate the sensor measurements and control actions as

$$y = \begin{bmatrix} y_1^T & \ldots & y_n^T \end{bmatrix}^T \qquad u = \begin{bmatrix} u_1^T & \ldots & u_n^T \end{bmatrix}^T$$

and then further amalgamate G and partition K as

$$G = \begin{bmatrix} G_{11} & \ldots & G_{1n} \\ \vdots & & \vdots \\ G_{n1} & \ldots & G_{nn} \end{bmatrix} \qquad K = \begin{bmatrix} K_{11} & \ldots & K_{1n} \\ \vdots & & \vdots \\ K_{n1} & \ldots & K_{nn} \end{bmatrix}$$

Whenever a single subscript is used, such as in some of the diagrams for the examples, it refers to a row from this partition, such as $G_i : \mathcal{U} \to Y_i$, or $K_i : \mathcal{Y} \to U_i$.

2.1 Propagation Delays

For any pair of nodes i and j we consider the subsystem $G_{ij} : \mathcal{U}_j \to \mathcal{Y}_i$, and define the propagation delay p_{ij} as the amount of time before a controller action at subsystem j can affect an output at subsystem i as such

$$p_{ij} = \text{Delay}(G_{ij}) \qquad \text{for all } i, j \in 1, \dots n$$

2.2 Transmission Delays

For any pair of subsystems k and l we define the (total) transmission delay t_{kl} as the minimum amount of time before the controller of subsystem k may use outputs from subsystem l. Given these constraints, we can define the overall subspace of admissible controllers S such that $K \in S$ if and only if

$$\text{Delay}(K_{kl}) \geq t_{kl} \qquad \text{for all } k, l \in 1, \dots n$$

In Section 3.2 we will break these total transmission delays out into a pure transmission delay, representing the time it takes to communicate the information from one subsystem to another, and a computational delay, representing the time it takes to process the information before it is used by the controller.

Triangle inequality For the main result of this paper, we will assume that the triangle inequality holds amongst the transmission delays, that is,

$$t_{ki} + t_{ij} \geq t_{kj} \qquad \text{for all } k, i, j$$

This is typically a very reasonable assumption for the following reasons. t_{kj} is defined as the minimum amount of time before controller k can use outputs from subsystem j. So if there existed an i such that the inequality above failed, that would mean that controller k could receive that information more quickly if it came indirectly via controller i. We would thus reroute this information to go through i, t_{kj} would be reset to $t_{ki} + t_{ij}$, and the inequality would hold.

To put it another way, we could think of each subsystem as a node on a directed graph, with the initial distance from any node j to any node k as t_{kj}, the time it takes before controller k can directly use outputs from subsystem j. We then want to find the shortest overall time for any controller k to use outputs from any subsystem j, that is, the shortest path from node j to node k. So to find our final t_{kj}'s, we run Bellman-Ford or another shortest path algorithm on our initial graph [2], and the resulting delays are thus guaranteed to satisfy the triangle inequality.

Note that the triangle inequality is never assumed to hold for the propagation delays.

2.3 Problem Formulation

Given a generalized plant P and transmission delays t_{kl} for each pair of subsystems, we define S as above, and we would then like to solve the following problem:

$$\begin{aligned} \text{minimize} \quad & \|f(P, K)\| \\ \text{subject to} \quad & K \text{ stabilizes } P \\ & K \in S \end{aligned} \qquad (1)$$

Here $\|\cdot\|$ is any norm on the closed-loop map chosen to encapsulate the control performance objectives. The delays associated with dynamics propagating from one subsystem to another are embedded in P. The subspace of admissible controllers, S, has been defined to encapsulate the constraints on how quickly information may be passed from one subsystem to another. We call the subspace S the ***information constraint***.

Many decentralized control problems may be expressed in the form of problem (1). In this paper, we focus on the case where S is defined by delay constraints as discussed above.

This problem is made substantially more difficult in general by the constraint that K lie in the subspace S. Without this constraint, the problem may be solved with many standard techniques. Note that the cost function $\|f(P,K)\|$ is in general a non-convex function of K. No computationally tractable approach is known for solving this problem for arbitrary P and S.

In Section 4 the transmission delays will be given between every pair of spatial points, rather than between every pair of subsystems, and the rest of the problem formulation remains the same.

2.4 Quadratic Invariance

In this subsection we define quadratic invariance, and give a brief overview of results regarding this condition, in particular, that it allows convex synthesis of optimal linear decentralized controllers.

Definition 1. *The set S is called **quadratically invariant** under G if*

$$KGK \in S \qquad \text{for all } K \in S$$

Note that, given G, we can define a quadratic map by $\Psi(K) = KGK$. Then a set S is quadratically invariant if and only if S is an invariant set of Ψ; that is, $\Psi(S) \subseteq S$.

It was shown in [7] that if S is a closed subspace and S is quadratically invariant under G, then with a change of variables, problem (1) is equivalent to the following optimization problem

$$\begin{aligned} \text{minimize} \quad & \|T_1 - T_2 Q T_3\| \\ \text{subject to} \quad & Q \in \mathcal{RH}_\infty \\ & Q \in S \end{aligned} \tag{2}$$

where T_1, T_2, T_3 are stable.

This is a convex optimisation problem. We may solve it to find the optimal Q, and then recover the optimal K for our original problem.

These results were achieved in [7] for operators acting on L_{2e} or ℓ_e, and similar results have been achieved [6] for other function spaces as well, also showing that quadratic invariance allows optimal linear decentralized control problems to be recast as convex optimization problems. These results can be directly applied to control over networks in Section 3. The technical work in [7] can be combined with the ideas in the Appendix of [1] to guarantee that this result holds for the operators considered in Section 4 as well; that is, for LTI spatio-temporal operators satisfying certain technical conditions.

The main focus of this paper is thus characterizing delays for which the information constraint S is quadratically invariant under the plant G.

3 Networks

In this section we consider control over networks, with propagation and transmission delays between nodes as defined and described above.

Section 3.1 contains the main result of the paper, where we prove that if this triangle inequality is satisfied, and if the propagation delay associated with any pair of subsystems is at least as large as the associated transmission delay, then the information constraint is quadratically invariant, and thus, optimal control may be cast as a convex optimization problem.

In Section 3.2 we break these total transmission delays out into a pure transmission delay, representing the time it takes to communicate the information from one subsystem to another, and a computational delay, representing the time it takes to process the information before it is used by the controller. We find, somewhat surprisingly, that transmitting faster than the propagation of dynamics still guarantees convexity, and in fact, that the computational delay causes the condition to be relaxed.

In Section 3.3, we discuss how sparsity constraints may be considered a special case of the framework analyzed in this paper, namely, by viewing them as very large delays. We then show how sparsity and delay constraints can be combined to handle the very general, realistic case of a network where some nodes are connected with delays and others are not connected at all.

We then consider a few examples in Section 3.4. First is an example corresponding to a very general problem of the control of vehicles in formation. The vehicles may have arbitrary positions, their dynamics propagate at a constant speed, and they communicate their measurements at a constant speed. The optimal control problem is amenable to convex synthesis as long as the communication speed exceeds the propagation speed. Even though this itself is a broad generalization of previously identified tractable classes, it follows almost immediately from the results of this paper.

Conditions are then derived for convexity of optimal control over a lattice, for two different types of assumptions on the propagation of dynamics.

3.1 Conditions for Convexity

We first provide a necessary and sufficient condition for quadratic invariance in terms of these delays, which is derived fairly directly from our definitions.

Theorem 2. *Suppose that G and S are defined as above. S is quadratically invariant under G if and only if*

$$t_{ki} + p_{ij} + t_{jl} \geq t_{kl} \qquad \text{for all } i,j,k,l \tag{3}$$

Proof. Given $K \in S$,

$$KGK \in S \iff \text{Delay}\big((KGK)_{kl}\big) \geq t_{kl} \text{ for all } k,l$$

We now seek conditions which cause this to hold.

$$(KGK)_{kl} = \sum_i \sum_j K_{ki} G_{ij} K_{jl}$$

and so for any k and l,

$$\begin{aligned}
&\text{Delay}\big((KGK)_{kl}\big) \\
&\qquad \geq \min_{i,j}\{\text{Delay}(K_{ki}G_{ij}K_{jl})\} \\
&\qquad \geq \min_{i,j}\{\text{Delay}(K_{ki}) + \text{Delay}(G_{ij}) + \text{Delay}(K_{jl})\} \\
&\qquad \geq \min_{i,j}\{t_{ki} + p_{ij} + t_{jl}\}
\end{aligned}$$

Thus S is quadratically invariant under G if

$$\min_{i,j}\{t_{ki} + p_{ij} + t_{jl}\} \geq t_{kl} \qquad \text{for all } k, l$$

which is equivalent to

$$t_{ki} + p_{ij} + t_{jl} \geq t_{kl} \qquad \text{for all } i, j, k, l$$

Now suppose that Condition (3) fails. Then there exists i, j, k, l such that

$$t_{ki} + p_{ij} + t_{jl} < t_{kl}$$

Consider K such that

$$K_{ab} = 0 \text{ if } (a,b) \notin \{(k,i),(j,l)\}$$

Then

$$(KGK)_{kl} \;=\; \sum_r \sum_s K_{kr} G_{rs} K_{sl} \;=\; K_{ki} G_{ij} K_{jl}$$

Since $\text{Delay}(G_{ij}) = p_{ij}$, we can easily choose K_{ki} and K_{jl} such that $\text{Delay}(K_{ki}) = t_{ki}$, $\text{Delay}(K_{jl}) = t_{jl}$, and

$$\text{Delay}\big((KGK)_{kl}\big) = t_{ki} + p_{ij} + t_{jl}$$

So $K \in S$ but $KGK \notin S$ and thus S is not quadratically invariant under G. ■

Main networks result. The following is the main result of this paper. It states that if the transmission delays satisfy the triangle inequality, and if the propagation delay between any pair of subsystems is at least as large as the transmission delay between those subsystems, then the information constraint is quadratically invariant. In other words, if along any link, data can be transmitted faster than dynamics propagate, then optimal controller synthesis may be cast as a convex optimization problem.

Theorem 3. *Suppose that G and S are defined as above, and that the transmission delays satisfy the triangle inequality. If*

$$p_{ij} \geq t_{ij} \qquad \textit{for all } i, j \tag{4}$$

then S is quadratically invariant under G.

Proof. Suppose Condition (4) holds. Then for all i, j, k, l we have

$$\begin{aligned} t_{ki} + p_{ij} + t_{jl} &\geq t_{ki} + t_{ij} + t_{jl} \\ &\geq t_{kl} \qquad \text{by the triangle inequality} \end{aligned}$$

and thus by Theorem 2, S is quadratically invariant under G. ■

Thus we have shown that the triangle inequality and Condition (4) are sufficient for quadratic invariance. The following remarks discuss assumptions under which they are necessary as well.

Remark 4. *If we assume that $t_{ii} = 0$ for all i, that is, that there is no delay before a subsystem's controller may use its own outputs, then we consider Condition* (3) *with $k = i$, $l = j$ and see that Condition* (4) *is necessary for quadratic invariance.*

Remark 5. *If we assume that $p_{ii} = 0$ for all i, that is, that there is no delay associated with propagating from a subsystem to itself, then we consider Condition* (3) *with $i = j$ and see that the triangle inequality is necessary for quadratic invariance.*

3.2 Computational Delays

In this subsection, we consider what happens when the controller of each subsystem has a computational delay c_i associated with it. The delay for controller i to use outputs from subsystem j, the total transmission delay, is then broken up into a pure transmission delay and this computational delay, as follows

$$t_{ij} = c_i + \tilde{t}_{ij}$$

If we were to assume that the triangle inequality held for the total transmission delays t_{ij} as before, then we would simply get the same results as in the previous section with the substitution above. In particular, we would find $p_{ij} \geq c_i + \tilde{t}_{ij}$ to be the condition for quadratic invariance. However, there are many cases where it makes sense to instead assume that the triangle inequality holds for the pure transmission delays $\tilde{t}_{ij}$, which is a stronger assumption. An example where such is clearly the case is provided in Section 3.4.1.

In this section we derive conditions for quadratic invariance when we can assume that the triangle inequality holds for the pure transmission delays $\tilde{t}_{ij}$, and get a surprising result.

As before, the propagation delays are defined as

$$p_{ij} = \text{Delay}(G_{ij}) \qquad \text{for all } i, j$$

and S is now defined such that $K \in S$ if and only if

$$\text{Delay}(K_{kl}) \geq c_k + \tilde{t}_{kl} \qquad \text{for all } k, l$$

Thus the necessary and sufficient condition for quadratic invariance from Theorem 2 becomes

$$c_k + \tilde{t}_{ki} + p_{ij} + c_j + \tilde{t}_{jl} \geq c_k + \tilde{t}_{kl} \qquad \text{for all } i, j, k, l$$

which reduces to

$$\tilde{t}_{ki} + p_{ij} + c_j + \tilde{t}_{jl} \geq \tilde{t}_{kl} \qquad \text{for all } i, j, k, l \tag{5}$$

The following theorem gives conditions under which the information constraint is quadratically invariant. It states that if the triangle inequality holds amongst the pure transmission delays, and if Condition (6) holds, then the information constraint is quadratically invariant. Surprisingly, we see that the computational delay now appears on the left side of the inequality. In other words, not only does transmitting data faster than dynamics propagate still allow for convex synthesis when we account for computational delay, but the condition is actually relaxed.

Theorem 6. *Suppose that G and S are defined as above, and that the pure transmission delays satisfy the triangle inequality. If*

$$p_{ij} + c_j \geq \tilde{t}_{ij} \qquad \textit{for all } i, j \tag{6}$$

then S is quadratically invariant under G.

Proof. Suppose Condition (6) holds. Then for all i, j, k, l we have

$$\begin{aligned}\tilde{t}_{ki} + p_{ij} + c_j + \tilde{t}_{jl} &\geq \tilde{t}_{ki} + \tilde{t}_{ij} + \tilde{t}_{jl} \\ &\geq \tilde{t}_{kl} \qquad \text{by the triangle inequality}\end{aligned}$$

and thus Condition (5) holds and S is quadratically invariant under G. ∎

Thus we have shown that the triangle inequality and Condition (6) are sufficient for quadratic invariance. The following remark discusses an assumption under which the condition is necessary as well.

Remark 7. *If we assume that $\tilde{t}_{ii} = 0$ for all i, that is, that there is no additional delay before a subsystem's controller may use its own outputs, other than the computational delay, then we consider Condition* (5) *with $k = i$, $l = j$ and see that Condition* (6) *is necessary for quadratic invariance. Since the computational delay has been extracted, this is now a very reasonable assumption which is essentially true by definition.*

3.3 Combining Sparsity and Delay Constraints

In this section, we discuss how sparsity constraints may be considered a special case of the framework analyzed in this paper. We then show how the two can be combined to handle the very general, realistic case of a network where some nodes are connected with delays as above and others are not connected at all. An explicit test for quadratic invariance in this case is provided.

The key observation is that a sparsity constraint may be considered an infinite delay. We thus define an extended notion of propagation and transmission delays, where they are assigned to be sufficiently large when they do not exist,and then the results from the rest of this paper may be applied to test for quadratic invariance and convexity.

3.3.1 Propagation Delays

We now consider a plant for which the controllers of certain subsystems may or may not have any effect on other subsystems, and when they do, there may be a propagation delay associated with that effect. First, define a binary matrix G^{bin} such that

$$G^{\text{bin}}_{ij} = \begin{cases} 0 & \text{if } G_{ij} = 0 \\ 1 & \text{otherwise} \end{cases}$$

In other words, G^{bin} defines the sparsity structure or interconnection structure of the plant, as $G^{\text{bin}}_{ij} = 0$ if subsystem i is not affected by inputs to subsystem j. We would then like to define the propagation delay p_{ij} to be extremely large if this is the case, as such

$$p_{ij} = \begin{cases} H & \text{if } G^{\text{bin}}_{ij} = 0 \\ \text{Delay}(G_{ij}) & \text{if } G^{\text{bin}}_{ij} = 1 \end{cases}$$

for some large H.

3.3.2 Transmission Delays

We similarly assign a binary matrix K^{bin} such that $K^{\text{bin}}_{kl} = 0$ if controller k may never use outputs from subsystem l. For any other pair of subsystems k and l we define the (total) transmission delay t_{kl} as in the rest of this paper; that is, as the minimum amount of time before the controller of subsystem k may use outputs from subsystem l. Given these constraints, we can define the overall subspace of admissible controllers S such that $K \in S$ if and only if

$$\begin{aligned} K_{kl} &= 0 && \text{for all } k, l \text{ such that } K^{\text{bin}}_{kl} = 0 \\ \text{Delay}(K_{kl}) &\geq t_{kl} && \text{for all } k, l \text{ such that } K^{\text{bin}}_{kl} = 1 \end{aligned}$$

We wish to assign a very large transmission delay to the former case, and so define

$$t_{kl} = H \qquad \text{for all } k, l \text{ such that } K^{\text{bin}}_{kl} = 0$$

for the same large H as above.

3.3.3 Condition for Convexity

Given these extended definitions of propagation delays and transmission delays for a combination of sparsity and delay constraints, we can now test for quadratic invariance using Theorem 2.

These definitions of extended delays along with our definition of the constraint set S allow us to use this and the rest of the results of this paper as long as H has been chosen large enough. Condition (3) is indeed necessary and sufficient for quadratic invariance as long as

$$H > 2\max\{t_{kl}\} + \max\{p_{ij}\}$$

where of course the first maximum is taken over all k, l such that $K^{\text{bin}}_{kl} = 1$ and the second is taken over all i, j such that $G^{\text{bin}}_{ij} = 1$. The bound on H arises because Condition (3) must fail if $K^{\text{bin}}_{kl} = 0$, but $K^{\text{bin}}_{ki} = G^{\text{bin}}_{ij} = K^{\text{bin}}_{jl} = 1$. It then also follows that the condition is satisfied if $K^{\text{bin}}_{ki} G^{\text{bin}}_{ij} K^{\text{bin}}_{jl} = 0$, as is also required, and that if all four of these values are 1, then whether the condition is satisfied depends on the delays, as in the rest of the paper.

3.4 Examples

We consider here some special cases of interest.

3.4.1 Vehicle Formation Example

We now consider an important special case, which corresponds to the problem of controlling multiple vehicles in a formation.

Suppose there are n subsystems (vehicles), with positions $x_1, \ldots, x_n \in \mathbb{R}^d$. Typically, we'll have $d = 3$, but these results hold for arbitrary $d \in \mathbb{Z}_+$.

Let D represent the diameter, the maximum distance between any two subsystems

$$D = \max_{i,j} \|x_i - x_j\|$$

For most applications of interest the appropriate norm throughout this section would be the Euclidean norm, but these results hold for arbitrary norm on $\mathbb{R}^d$.

We suppose that dynamics of all vehicles propagate at a constant speed, determined by the medium, such that the propagation delays are proportional to the distance between vehicles, as illustrated in Figure 3.

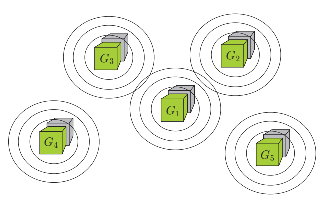

Figure 3: Communication and propagation in all directions

Let γ_p be the amount of time it takes dynamics to propagate one unit of distance, i.e., the inverse of the speed of propagation. For example, when considering formations of aerial vehicles, γ_p would equal the inverse of the speed of sound.

The system G is then such that

$$\text{Delay}(G_{ij}) = \gamma_p \|x_i - x_j\| \qquad \text{for all } i, j$$

We similarly suppose that data can be transmitted at a constant speed, such that the transmission delays are proportional to the distances between vehicles, such as if each vehicle could broadcast its information to the others. This is also illustrated in Figure 3. We assume that the perturbations from our desired formation are small enough that, for the purposes of controller synthesis, we may consider these delays to be fixed.

Let γ_t be the amount of time it takes to transmit one unit of distance, i.e., the inverse of the speed of transmission. Let C be the computational delay at each vehicle. The set of admissible controllers is then defined such that $K \in S$ if and only if

$$\text{Delay}(K_{kl}) \geq C + \gamma_t \|x_k - x_l\| \qquad \text{for all } k, l$$

We can now apply Theorem 6 with

$$p_{ij} = \gamma_p \|x_i - x_j\|, \ \tilde{t}_{ij} = \gamma_t \|x_i - x_j\|, \text{ and } c_i = C \qquad \text{for all } i, j$$

Clearly, $\tilde{t}_{ii} = 0$ for all i as in Remark 7, so the conditions of Theorem 6 are both necessary and sufficient for quadratic invariance.

Theorem 8. *Suppose that G and S are defined as above. S is quadratically invariant under G if and only if*

$$\gamma_p + (C/D) \geq \gamma_t$$

Proof. Since any norm satisfies the triangle inequality, the pure transmission delays clearly satisfy the triangle inequality, so applying Theorem 6, S is quadratically invariant under G if and only if

$$\gamma_p \|x_i - x_j\| + C \geq \gamma_t \|x_i - x_j\| \qquad \text{for all } i, j$$

which is equivalent to

$$\gamma_p + (C/D) \geq \gamma_t$$

■

Thus we see that, in the absence of computational delay, finding the minimum-norm controller may be reduced to a convex optimization problem when the speed of transmission is faster than the speed of propagation; that is, when $\gamma_p \geq \gamma_t$. We also see that this not only remains true in the presence of computational delay, but that we get a buffer relaxing the condition.

A similar result was previously achieved for a very specific case of vehicles equally spaced along a line [7]. This shows how the results of this paper allow us to effortlessly generalize to the case considered in this subsection, where the vehicles have arbitrary positions in arbitrary dimensions. This is a crucial generalization for applications to realistic formation flight problems.

3.4.2 Two-Dimensional Lattice Example

In this subsection we will consider subsystems distributed in a lattice, and use these results to derive the conditions for convexity of the associated optimal decentralized control problem.

We first consider the case where the controllers can communicate along the edges of the lattice with a delay of t, and the dynamics similarly propagate along the edges with a delay of p, as illustrated in Figure 4.

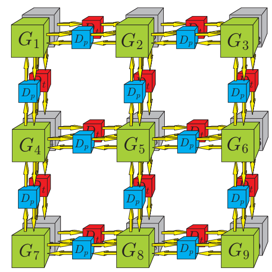

Figure 4: Two-dimensional lattice with dynamics propagating along edges

It is a straightforward consequence of this paper that the optimal controllers may be synthesized with convex programming if

$$p \geq t$$

We now consider a more interesting variant, where the controllers again communicate only along the edges of the lattice, but now the dynamics propagate in all directions, as illustrated in Figure 5.

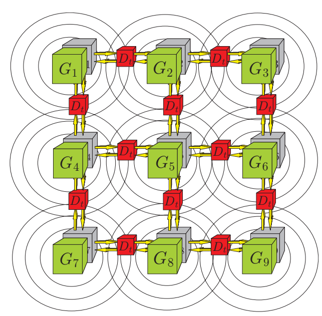

Figure 5: Two-dimensional lattice with dynamics propagating in all directions

Let p again be the amount of time it takes for the dynamics to propagate one edge length, for instance, from G_1 to G_4. Along a diagonal then, for example, between G_1 and G_5, the propagation delay is $p\sqrt{2}$ and the transmission delay is $2t$. The condition for convexity therefore becomes

$$p \geq t\sqrt{2}$$

4 Spatio-Temporal Systems

In this section we turn our attention to spatio-temporal systems. We consider delays before control actions at certain points in the spatial domain can effect outputs at other points, as well as delays before control actions at certain points may depend on outputs at other points. These are analogous to the propagation delays and transmission delays defined between nodes/subsystems in the previous section.

In Section 4.1, we define these systems, these delays, and the subsequent information constraints. This section is a very technical necessity; the main message for the reader wishing to skip it is that we can define the delay between two points for a spatio-temporal system, and then build up a framework analogous to that for control over networks.

We then show in Section 4.2 that once the framework is established in a particular way, we can derive results for the convexity of the optimal distributed control problem which are a clear extension to the results for networks; that is, if the transmission delays satisfy the triangle inequality, then the information constraint is quadratically invariant if the transmission delay between any two points is less than the propagation delay between those points.

We then in Section 4.3 consider the special case of spatially invariant systems. It falls out almost immedaitely from our more general spatio-temporal result that these problems are convex if the propagation function is subadditive, which is itself a broad generalisation of previously identified convex problems of this class.

These results hold for either continuous or discrete spatial or temporal domains. When the spatial domain is discrete, the results and the proofs are virtually identical to those of the previous section, so we will focus here on the continuous spatial domain.

4.1 Spatio-temporal preliminaries

Suppose we have d spatial dimensions for some $d \in \mathbb{Z}_{>0}$. We consider systems which are linear and time-invariant (though not necessarily spatially invariant) and that we can thus characterize in terms of a spatio-temporal impulse response $h : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}$. We can then define a system H_h as $H_h u = y$ where

$$y(x,t) = \int_{\mathbb{R}^d} \int_{\mathbb{R}} h(x-\xi, t-\tau) u(\xi,\tau)\, d\tau\, d\xi \qquad \text{for } x \in \mathbb{R}^d,\ T \in \mathbb{R}$$

We further restrict ourselves to systems for which the impulse response satisfies

$$\sup_{x \in \mathbb{R}^d} \sup_{\tau \in [0,T]} \int_{\mathbb{R}^d} |h(x,\xi,\tau)|\, d\xi < \infty \qquad \text{for all } T \in \mathbb{R}_+$$

and

$$\sup_{\xi \in \mathbb{R}^d} \sup_{\tau \in [0,T]} \int_{\mathbb{R}^d} |h(x,\xi,\tau)|\, dx < \infty \qquad \text{for all } T \in \mathbb{R}_+$$

and (temporal) causality requires that $h(x,\xi,\tau) = 0$ for all $\tau < 0$.

Delays. We now seek to define $\text{Delay}(H_h, x, \xi)$ as the delay for such an operator from point $\xi \in \mathbb{R}^d$ to point $x \in \mathbb{R}^d$, analagous to how in the previous section $\text{Delay}(G_{ij})$ gave the delay from node j to node i. However, we must exercise great care with this definition. If we were to similarly define it as the smallest time until a change in (and only in) $u(\xi,\cdot)$ could affect a change in $y(x,\cdot)$, then this would always be infinite,, since only changing the input on a set of measure zero cannot alter the output.

Similarly, if we defined it directly in terms of a given impulse response as the smallest time for which $h(x,\xi,\tau) \neq 0$, this would make distinctions that we do not want to make, since two impulse responses that differ only on a set of measure zero will yield the same mapping, but would then give different delay profiles.

We thus define the delay as follows for all $x, \xi \in \mathbb{R}^d$

$$\text{Delay}(H_h, x, \chi) = \sup_{\varepsilon > 0} \operatorname{ess\,inf} \{ \tau \geq 0 \mid h(x',\xi',\tau) \neq 0, \|(x',\xi') - (x,\xi)\| < \varepsilon \}$$

This gives the same delay profile for any impulse responses which differ on only a set of measure zero, and thus for any which characterise a given mapping. Further, letting $t(x,\xi) = \text{Delay}(H,x,\xi)$ for all $x,\xi \in \mathbb{R}^d$, this will always give a lower semicontinuous function. Such a function has a closed epigraph, which will be essential for our purposes.

If we started with an impulse response that had its support in the epigraph of a lower semicontinuous function, then that function would result as the delay profile, and the delay function would revert to the more intuitive

$$\inf\{\tau \geq 0 \mid h(x, \xi, \tau) \neq 0\}$$

Lastly, this results in the following inequalities for the delays of a composition or an addition of operators:

$$\text{Delay}(AB, \xi_k, \xi_j) \geq \inf_{\xi_i} \text{Delay}(A, \xi_k, \xi_i) + \text{Delay}(B, \xi_i, \xi_j)$$
$$\text{Delay}(A + B, \xi_k, \xi_j) \geq \min\{\text{Delay}(A, \xi_k, \xi_j), \text{Delay}(B, \xi_k, \xi_j)\}$$

Propagation delays. Given a plant G, we define a propagation delay function giving the amount of time before a controller action around point ξ_j can affect an output around point ξ_i as

$$p(\xi_i, \xi_j) = \text{Delay}(G, \xi_i, \xi_j) \qquad \text{for all } \xi_i, \xi_j \in \mathbb{R}^d$$

Transmission delays. For any pair of points $\xi_j, \xi_l \in \mathbb{R}^d$ we define the transmission delay $t(\xi_k, \xi_l)$ as the minimum amount of time before the controller in a neighbourhood of point ξ_k may use outputs from a neighbourhood of point ξ_l. Given these constraints, we can define the overall subspace of admissible controllers S such that $K \in S$ if and only if

$$\text{Delay}(K, \xi_k, \xi_l) \geq t(\xi_k, \xi_l) \qquad \text{for all } \xi_k, \xi_l \in \mathbb{R}^d$$

If t is lower semicontinuous, then the information constraint S can also be characterised by $K \in S$ if and only if K can be represented by a kernel h_K such that

$$h_K(\xi_k, \xi_l, \tau) = 0 \qquad \text{for all } \tau < t(\xi_k, \xi_l)$$

Triangle inequality. For the main spatio-temporal result, we will again assume that the triangle inequality holds amongst the transmission delays, which is written as

$$t(\xi_k, \xi_i) + t(\xi_i, \xi_j) \geq t(\xi_k, \xi_j) \qquad \text{for all } \xi_k, \xi_i, \xi_j \in \mathbb{R}^d$$

and which is again a reasonable assumption for the same reasons as in the case of networks.

4.2 Conditions for convexity

We are now ready to prove our main spatio-temporal results.

We first provide a necessary and sufficient condition for quadratic invariance in terms of these delays, which is again derived fairly directly from our definitions.

The proof of this, as well as the proof of the main result of this section, follow very similarly to those for networks.

Theorem 9. *Suppose that G and S are defined as above. S is quadratically invariant under G if and only if for all $\xi_i, \xi_j, \xi_k, \xi_l \in \mathbb{R}^d$*

$$t(\xi_k, \xi_i) + p(\xi_i, \xi_j) + t(\xi_j, \xi_l) \geq t(\xi_k, \xi_l) \tag{7}$$

Proof. Given $K \in S$,

$$KGK \in S \iff \text{Delay}(KGK, \xi_k, \xi_l) \geq t(\xi_k, \xi_l) \qquad \text{for all } \xi_k, \xi_l \in \mathbb{R}^d$$

We now seek conditions which cause this to hold. For all $\xi_k, \xi_l \in \mathbb{R}^d$,

$$\text{Delay}(KGK, \xi_k, \xi_l) \geq \inf_{\xi_i, \xi_j} \{t(\xi_k, \xi_i) + p(\xi_i, \xi_j) + t(\xi_j, \xi_l)\}$$

Thus S is quadratically invariant under G if

$$\inf_{\xi_i, \xi_j} \{t(\xi_k, \xi_i) + p(\xi_i, \xi_j) + t(\xi_j, \xi_l)\} \geq t(\xi_k, \xi_l) \qquad \text{for all } \xi_k, \xi_l$$

which is equivalent to

$$t(\xi_k, \xi_i) + p(\xi_i, \xi_j) + t(\xi_j, \xi_l) \geq t(\xi_k, \xi_l) \qquad \text{for all } \xi_i, \xi_j, \xi_k, \xi_l \in \mathbb{R}^d$$

Now suppose that Condition (7) fails. Then there exists $\xi_i, \xi_j, \xi_k, \xi_l \in \mathbb{R}^d$ such that

$$t(\xi_k, \xi_i) + p(\xi_i, \xi_j) + t(\xi_j, \xi_l) < t(\xi_k, \xi_l)$$

Consider K such that

$$h_K(\xi_a, \xi_b, \tau) = \begin{cases} \gamma & \text{if } \|(\xi_a, \xi_b) - (\xi_j, \xi_l)\| < \varepsilon,\ \tau \in [t(\xi_a, \xi_b), t(\xi_a, \xi_b) + \beta] \\ \gamma & \text{if } \|(\xi_a, \xi_b) - (\xi_k, \xi_i)\| < \varepsilon,\ \tau \in [t(\xi_a, \xi_b), t(\xi_a, \xi_b) + \beta] \\ 0 & \text{otherwise} \end{cases}$$

for some $\beta > 0$, $\gamma \neq 0$, $\varepsilon > 0$. Then

$$\text{Delay}(KGK, \xi_k, \xi_l) \leq t(\xi_k, \xi_i) + p(\xi_i, \xi_j) + t(\xi_j, \xi_l)$$

So $K \in S$ but $KGK \notin S$ and thus S is not quadratically invariant under G. ∎

Main spatio-temporal result. The following is the main spatio-temporal result of this paper. It states that if the transmission delays satisfy the triangle inequality, and if the propagation delay between any pair of spatial points is at least as large as the transmission delay between those points, then the information constraint is quadratically invariant. In other words, if between any points, data can be transmitted faster than dynamics propagate, then optimal controller synthesis may be cast as a convex optimisation problem.

Theorem 10. *Suppose that G and S are defined as above, and that the transmission delays satisfy the triangle inequality. If*

$$p(\xi_i, \xi_j) \geq t(\xi_i, \xi_j) \qquad \textit{for all } \xi_i, \xi_j \tag{8}$$

then S is quadratically invariant under G.

Proof. Suppose Condition (8) holds. Then for all $\xi_i, \xi_j, \xi_k, \xi_l$ we have

$$\begin{aligned} t(\xi_k, \xi_i) + p(\xi_i, \xi_j) + t(\xi_j, \xi_l) &\geq t(\xi_k, \xi_i) + t(\xi_i, \xi_j) + t(\xi_j, \xi_l) \\ &\geq t(\xi_k, \xi_l) \qquad \text{by the triangle inequality} \end{aligned}$$

and thus by Theorem 3, S is quadratically invariant under G. ∎

Thus we have shown that the triangle inequality and Condition (8) are sufficient for quadratic invariance. The following remarks discuss assumptions under which they are necessary as well.

Remark 11. *If we assume that $t(\xi_i, \xi_i) = 0$ for all $\xi_i \in \mathbb{R}^d$, that is, that there is no delay before a controller at a given point may use the output from the same point, then we consider Condition* (7) *with $\xi_k = \xi_i$, $\xi_l = \xi_j$ and see that Condition* (8) *is necessary for quadratic invariance.*

Remark 12. *If we assume that $p(\xi_i, \xi_i) = 0$ for all $\xi_i \in \mathbb{R}^d$, that is, that there is no delay before a control action at a given point can affect the output at that point, then we consider Condition* (7) *with $\xi_i = \xi_j$ and see that the triangle inequality is necessary for quadratic invariance.*

4.3 Spatial invariance

We now consider the special case of spatial invariance, where the behaviour of a spatio-temporal system from one point to another remains constant if both points undergo the same shift; in other words, we can express the impulse response as

$$h(\xi_i, \xi_j, \tau) = h^{\mathrm{SI}}(\xi_i - \xi_j, \tau) \qquad \text{for all } \xi_i, \xi_j \in \mathbb{R}^d,\ \tau \in \mathbb{R}$$

for a spatially invariant impulse response $h^{\mathrm{SI}} : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}$. (We hereafter drop the superscript and assume that we are dealing with spatially invariant impulse responses which are thus functions of one spatial variable and one temporal variable.) We can then define a system H_h as $H_h u = y$ where

$$y(x, T) = \int_{\mathbb{R}^d} \int_{\mathbb{R}} h(x - \xi, T - \tau) u(\xi, \tau)\, d\tau\, d\xi$$

Our previous restrictions now imply that the impulse response satisfies

$$\sup_{\tau \in [0,T]} \int_{\mathbb{R}^d} |h(x, \tau)|\, dx < \infty \qquad \text{for all } T \in \mathbb{R}_+$$

and (temporal) causality requires that $h(x, \tau) = 0$ for all $\tau < 0$, as in [1].

This will result in a spatially invariant propagation delay function $p^{\mathrm{SI}} : \mathbb{R}^d \to \mathbb{R}_+$ given by

$$p(\xi_i, \xi_j) = p^{\mathrm{SI}}(\xi_i - \xi_j) \qquad \text{for all } \xi_i, \xi_j \in \mathbb{R}^d$$

and we hereafter drop this superscript as well.

We now similarly consider the special case of imposing spatially invariant constraints on our controller. Then we could express the transmission delays which determine our constraint set S_t as

$$t(\xi_k, \xi_l) = f(\xi_k - \xi_l) \qquad \text{for all } \xi_k, \xi_l \in \mathbb{R}^d$$

for some function $f : \mathbb{R}^d \to \mathbb{R}$, which we then, following the notation of [1], call a propagation function. The set of admissible controllers S_f, defined such that an input in a neighbourhood of location ξ cannot affect the output in a neighbourhood of location $x+\xi$ until at least $f(x)$ units of time have elapsed, then corresponds to our previous definition of S_t with the above equality. (Note that this is a slight break from the notation of [1].) More formally,

$$S_f = \{K \mid \mathrm{Delay}(K, \xi_k, \xi_l) \geq f(\xi_k - \xi_l) \qquad \text{for all } \xi_k, \xi_l \in \mathbb{R}^d\}$$

which, if f is lower semicontinuous, is equivalent to

$$\{H_h \mid h(x, \tau) = 0 \text{ for all } \tau < f(x)\}$$

Note that restricting these sets to only contain causal operators now corresponds to only considering propagation functions for which $f(x) \geq 0$ for all $x \in R^d$.

In [1], focusing on $d = 1$, it was shown that finding the optimal controller constrained to such a set ($K \in S_f$) could be cast as a convex optimisation problem if

1. $G \in S_f$,
2. $f(0) = 0$,
3. $f(x)$ is concave in $\mathbb{R}_+$,
4. $f(x)$ is concave in $\mathbb{R}_-$.

Such a function $f(\cdot)$ is shown in Figure 6, and we see the motivation for calling these systems funnel causal.

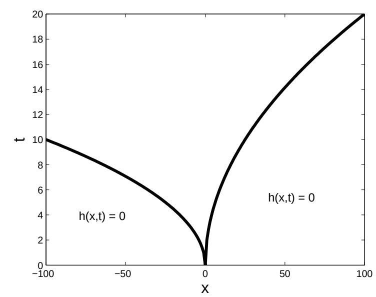

Figure 6: Funnel causality

Main spatially invariant result. We now turn the results of this paper onto this set of problems to see if we can generalize the class which is amenable to convex synthesis.

We show that if the plant is slower than the controller, the propagation function only needs to be *subadditive* for the information contraint to be quadratically invariant.

Theorem 13. *Suppose G and S_f are defined as above. Then S_f is quadratically invariant under G if the following hold*

(i) $p(x) \geq f(x)$ *for all* $x \in \mathbb{R}^d$ *(i.e.,* $G \in S_f$*)*

(ii) $f(x+y) \leq f(x) + f(y)$ *for all* $x, y \in \mathbb{R}^d$

Proof. The condition that the dynamics cannot propagate from one point to another any faster than the controllers are allowed to communicate (8) is indeed equivalent to Condition (i) above. It remains to see how the triangle inequality amongst transmission delays manifests itself in this framework.

This remaining condition (triangle inequality) for quadratic invariance from Theorem 10

$$t(\xi_k, \xi_j) \leq t(\xi_k, \xi_i) + t(\xi_i, \xi_j) \qquad \text{for all } \xi_k, \xi_i, \xi_j \in \mathbb{R}^d$$

reduces with spatial invariance to

$$f(\xi_k - \xi_j) \leq f(\xi_k - \xi_i) + f(\xi_i - \xi_j) \qquad \text{for all } \xi_k, \xi_i, \xi_j \in \mathbb{R}^d$$

which, since we can choose $x = \xi_k - \xi_i$, $y = \xi_i - \xi_j$, or conversely, $\xi_j = 0$, $\xi_i = y$, $\xi_k = x+y$, is equivalent to

$$f(x+y) \leq f(x) + f(y) \qquad \text{for all } x, y \in \mathbb{R}^d$$

■

When $d = 1$, this includes, but is not limited to, functions which satisfy Conditions 2-4.

4.3.1 Spatially invariant examples

We now consider the case of one spatial dimension to explore some constraint sets which are not funnel causal but which we now know to also yield convex optimization problems. One type of function which is subadditive but not concave is the step function, $f(x) = \frac{\lceil x \rceil}{B}$. This is shown is Figure 7, with different values of B for positive and negative x.

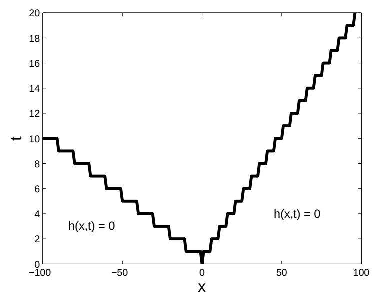

Figure 7: Subadditive step function

Note that it is important that the ceiling function be used rather than the floor, both to ensure that the function is subadditive (and thus, that the associated information constraint is quadratically invariant) as well as to ensure that the function is lower semi-continuous (and thus, that the epigraph and associated information constraint are closed).

While not explicitly stated, an implicit constraint in Conditions 2-4, when combined with the nonnegativity of f, is that f has to be nonincreasing in $\mathbb{R}_-$ and nondecreasing in $\mathbb{R}_+$. This would typically be desired, as it corresponds to saying that information can be transmitted at least as fast over a shorter distance as it can over a longer distance. However, it is interesting to note that this is no longer required either. An example of a subadditive function which violates this assumption is shown in Figure 8.

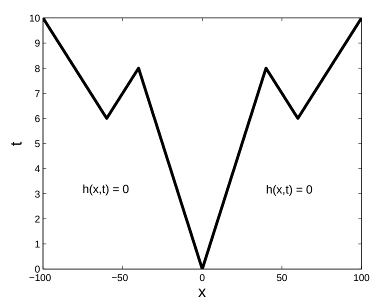

Figure 8: Subadditive non-monotonic function

We mostly focused on continuous spatial domains in this paper, but noted that all of the results hold for discrete spatial domains as well. The propagation function shown in Figure 9,

$$f(x) = \begin{cases} 0 & \text{if } x \text{ even} \\ 1 & \text{if } x \text{ odd} \end{cases}$$

is an example of a propagation function on $\mathbb{Z}$ which is indeed subadditive, but which is not at all concave or monotonic.

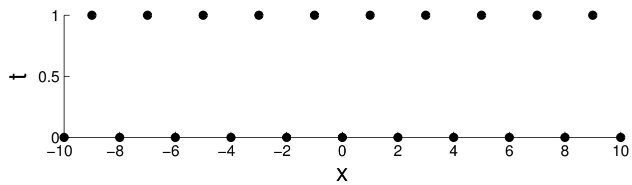

Figure 9: Subadditive discrete function

Two other generalizations have been achieved which are more subtle but likely more important. The first is that Condition 2 is no longer required, and we can have any $f(0) \geq 0$, which could be essential for incorporating computational delay. The second is that, in this framework, it was seamless to consider multiple spatial dimensions d.

5 Conclusions

We have studied the problem of finding optimal controllers subject to delays between different parts of the controller.

We first studied this in the context of control over networks, where multiple subsystems are subject to constraints on how quickly they can share information In Theorem 3 we showed that, presuming the transmission delays satisfy the triangle inequality, if the transmission delay between any pair of subsystems is less than the corresponding propagation delay, then the information constraint is quadratically invariant. This allows for convex synthesis of the optimal decentralized controllers.

We then studied this problem in the context of spatio-temporal systems, subject to constraints on how quickly information can be shared between given points. We showed in Theorem 10 that a result analogous to that for networks exists in this context, and that the optimal distributed control problem is convex when the transmission delay between any pair of points is less than the corresponding propagation delay, provided that a similar triangle inequality holds.

We further showed in Theorem 13 that if we apply this result to the special case of spatial invariance, it allows us to broadly extend the class of such systems amenable to convex controller synthesis, to systems in an arbitrary number of dimensions supported by any subadditive propagation function.

Delay constraints which allow for convex controller synthesis have thus been simply characterised, unified, and broadly generalised.

Acknowledgements

The authors would like to thank Michael Cantoni for several helpful discussions.

References

[1] B. Bamieh and P. G. Voulgaris. A convex characterization of distributed control problems in spatially invariant systems with communications constraints. *Systems and Control Letters*, 54(6):575–583, June 2005.

[2] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization*. Prentice-Hall, Englewood Cliffs, NJ, 1982.

[3] X. Qi, M. Salapaka, P.G. Voulgaris, and M. Khammash. Structured optimal and robust control with multiple criteria: A convex solution. *IEEE Transactions on Automatic Control*, 49(10):1623–1640, 2004.

[4] M. Rotkowitz. Parameterization of causal stabilizing controllers over networks with delays. In *Proc. IEEE Global Telecommunications Conference*, pages 1–5, November 2006.

[5] M. Rotkowitz, R. Cogill, and S. Lall. A simple condition for the convexity of optimal control over networks with delays. In *Proc. IEEE Conference on Decision and Control*, pages 6686–6691, 2005.

[6] M. Rotkowitz and S. Lall. Decentralized control information structures preserved under feedback. In *Proc. IEEE Conference on Decision and Control*, pages 569–575, December 2002.

[7] M. Rotkowitz and S. Lall. A characterization of convex problems in decentralized control. *IEEE Transactions on Automatic Control*, 51(2):274–286, February 2006.

[8] P. G. Voulgaris. Control under structural constraints: An input-output approach. In *Lecture notes in control and information sciences*, pages 287–305, 1999.

[9] H. S. Witsenhausen. Separation of estimation and control for discrete time systems. *Proceedings of the IEEE*, 59(11):1557–1566, 1971.

Biographical notes

Michael Rotkowitz is the Future Generation Fellow in the Department of Electrical and Electronic Engineering at the University of Melbourne, as well as an Honorary Fellow in the Department of Mathematics and Statistics. He received his BS in Mathematical and Computational Science, MS in Statistics, and MS and PhD in Aeronautics and Astronautics, all from Stanford University. He also worked for J.P. Morgan Investment Management, New York, 1996–1998. He was a Postdoctoral fellow at the Royal Institute of Technology (KTH), 2005–2006, and at the Australian National University, 2006–2008. His awards include the IEEE CDC Best Student Paper Award, the IFAC World Congress Young Author Prize, and the George S. Axelby Outstanding Paper Award.

Randy Cogill earned his PhD in Electrical Engineering from Stanford University in 2007. Since August 2007, he has been an Assistant Professor in the Department of Systems and Information Engineering at University of Virginia. His research interests fall into the broad categories of control, optimisation, applied stochastic processes, and networks. He is particularly interested in problems arising in communication networks, and the effects of scheduling and coding strategies on network dynamics. He was a recipient of the Stanford Graduate fellowship, and was co-awarded the best student paper award at the 2005 IEEE Conference on Decision and Control.

Sanjay Lall is an Associate Professor of Electrical Engineering, Associate Professor of Aeronautics and Astronautics, and Vance D. and Arlene C. Coffman Faculty Scholar at Stanford University. Until 2000, he was a Research fellow at the California Institute of Technology, and prior to that he was NATO Research Fellow at Massachusetts Institute of Technology. He received the PhD in Engineering and BA in Mathematics from the University of Cambridge, England. He received the George S. Axelby Outstanding Paper Award by the IEEE Control Systems Society in 2007, the NSF Career award in 2007, Presidential Early Career Award for Scientists and Engineers (PECASE) in 2007, and the Graduate Service Recognition Award from Stanford university in 2005.