

Performance Certification of Interconnected Systems using Decomposition Techniques

Chris Meissen¹

Laurent Lessard¹

Andrew Packard¹

American Control Conference, pp. 5030–5036, 2014

Abstract

We propose a unified framework for the certification of stability and input-output performance of interconnected dynamical systems. In our approach, we seek local dissipativity certificates for each subsystem such that when they are combined, the performance of the entire interconnected system is certified. We also demonstrate, by the use of numerical simulations, that the Alternating Direction Method of Multipliers (ADMM) is a promising computational approach for solving such problems.

1 Introduction

Many systems of interest naturally occur as an interconnection of a large number of smaller subsystems. Examples include: power systems, biological networks, autonomous agents, and communication networks. It is often the case that the individual components are well-modeled and characterized, but the global interconnected system is prohibitively difficult to analyze. The reason for this is that the computational tools used to analyze and certify the performance of the small subsystems do not scale well to larger systems. In this paper, we consider a general interconnection of subsystems as depicted in Figure 1.




Figure 1: Interconnected system with input and output.

The given subsystems are G_i which map $u_i \mapsto y_i$, and M is a static matrix that characterizes the interconnection topology. Although we think of M as being sparse, our approach does not rely on this fact and is applicable for any dense M .

We consider performance in the dissipativity framework [20]; specifically the case with quadratic supply rates [21]. The choice of supply rate dictates the property that is to be verified. For example, we may ask questions such as: “is the interconnected system stable?”, “is the L_2 gain from d to e less than 2?”, or “is the map from d to e passive?”. Given a supply rate, the goal is to find an associated storage function, which may be thought of as a generalized Lyapunov function. If a storage function is found, then the property is verified and the storage function serves as a certificate.

In general, when the G_i are nonlinear systems, determining dissipativity is very difficult. In the case of linear, time-invariant systems, the problem can be formulated as a linear matrix inequality and can be solved reliably [21]. For nonlinear, time-invariant systems with polynomial dynamics, a sum-of-squares relaxation can be used [17, 19]. The computational complexity of this latter approach grows dramatically with both the state dimension and the order of the polynomial certificates.

Main contributions. This paper has two main contributions. The first is to pose the dissipativity-verification question as a distributed optimization problem. Roughly, we search over arbitrary dissipativity properties satisfied by the local subsystems in such a way that the desired dissipativity property for the interconnected system emerges. This scalable formulation both generalizes recent efforts to characterize the performance of interconnected systems, and paves the way for further generalizations to more complicated performance criteria.

Our second contribution is to argue that ADMM, a popular large-scale optimization algorithm [6], is a potentially viable approach for solving our proposed formulation. We support our claim with numerical evidence, by solving randomly generated interconnections of 50 linear systems and comparing various optimization algorithms.

The paper is organized as follows. We first review relevant prior work and required preliminaries in Section 2. In Section 3, we formulate the distributed optimization problem, in Section 4 we discuss various decomposition methods including ADMM, and in Section 5 we present the results of our numerical simulations. We conclude in Section 6 with a discussion of possible generalizations.

¹C. Meissen, L. Lessard, and A. Packard are with the Department of Mechanical Engineering at the University of California, Berkeley, CA 94720, USA.
{cmeissen, lessard, apackard}@berkeley.edu

Prior work. Compositional analysis of large-scale interconnected dynamical systems has been studied extensively. The conventional approach, as presented in [2, 3, 14, 20] and others is to fix supply rates and storage functions for each of the local subsystems. Then, a candidate for a global storage function is sought as a linear combination of the local storage functions. The method presented in this paper is far less conservative in that we also optimize over arbitrary local storage functions and supply rates. Thus, global performance may be certified via local certificates that have been automatically generated.

The idea of optimizing over the local supply rates and storage functions was first introduced in [18]. The present work generalizes the work in [18] in two major ways. First, we consider dissipativity with respect to a quadratic supply rate rather than only considering stability certification. Second, we again consider dissipativity with respect to a general quadratic supply rate for the local subsystems, whereas in [18] the supply rates were constrained to be diagonally-scaled induced L_2 -norms.

While we assume the system has a natural decomposition into simpler subsystems, this may not always be the case. Recent works have explored the idea of automatically finding favorable decompositions for such problems. In [1, 2], a decomposition algorithm based on dissipation inequalities was proposed. Another approach is a Hankel-norm based lumping technique [13]. These approaches may be used to partition a large-scale system into the canonical form of Figure 1.

2 Preliminaries

Dissipativity theory. See [20] for a more in-depth review of this topic. Consider a time-invariant, continuous-time dynamical system described by

$$\begin{aligned} \dot{x}(t) &= f(x(t), u(t)), & f(0, 0) &= 0 \\ y(t) &= h(x(t), u(t)), & h(0, 0) &= 0 \end{aligned} \quad (1)$$

with $x(t) \in \mathbb{R}^n$, $u(t) \in \mathbb{R}^m$, and $y(t) \in \mathbb{R}^p$. A *supply rate* is a function $w : \mathbb{R}^m \times \mathbb{R}^p \rightarrow \mathbb{R}$. A system of the form (1) is said to be *dissipative* with respect to a supply rate w if there exists a differentiable and nonnegative function $V : \mathbb{R}^n \rightarrow \mathbb{R}_+$ such that

$$\nabla V(x)^\top f(x, u) - w(u, y) \leq 0 \quad (2)$$

for all $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, and $y = h(x, u)$. Equation (2) is referred to as the *Dissipation Inequality Equation* (DIE) and V as a *storage function*. For a given supply rate, the storage function is generally not unique if it exists.

In the special case where the system is autonomous and the supply rate w is equal to zero, if V is positive definite then (2) certifies stability of the system and V is referred to as a *Lyapunov function*.

If f and h are linear and the supply rate w is a quadratic function, then the system is dissipative if and only if there exists a quadratic storage function V . This problem can then be cast as a linear matrix inequality (LMI). If f and h have polynomial dynamics and the storage function V is restricted to be a polynomial function of fixed order then the problem of certifying dissipativity can be relaxed to a sum-of-squares (SOS) program.

A key idea is that finding a storage function V is difficult. Much like searching for Lyapunov functions, one must often broaden the search to very complicated classes of functions before a suitable one is found.

Interconnected systems. Consider Figure 1. The system consists of N known subsystems, G_i , with a known, static interconnection $M \in \mathbb{R}^{m \times p}$. Therefore,

$$\begin{bmatrix} u \\ e \end{bmatrix} = M \begin{bmatrix} y \\ d \end{bmatrix} \quad (3)$$

Each G_i has dynamics of the form (1) and is characterized by a local state x_i and output functions f_i and g_i , respectively. We assume that the interconnected system is well-posed, meaning that for any $d \in L_{2e}$, and any initial condition x_0 there exists unique $e, u, y \in L_{2e}$ that causally depend on d .

We would like to know whether the interconnected system is dissipative with respect to some given supply rate w_0 . Our approach is twofold.

1. For each i , find a local supply rate w_i such that G_i is dissipative with respect to w_i . This means that there exists some storage function V_i such that

$$\nabla V_i(x_i)^\top f_i(x_i, u_i) - w_i(u_i, y_i) \leq 0 \quad (4)$$

holds for all x_i and u_i , with $y_i = h_i(x_i, u_i)$.

2. Check whether the interconnected system is dissipative with respect to w_0 with storage function that is the sum of the local storage functions. So if we define $V(x) = \sum_{i=1}^N V_i(x_i)$, then

$$\nabla V(x)^\top f(x, u) - w_0(d, e) \leq 0 \quad (5)$$

holds for all x and d , with $y = h(x, u)$ and (3).

In our approach, (4)–(5) are combined to eliminate the storage functions V_i , as we will see in Section 3. Therefore, verifying step 2 is carried out without explicit knowledge of the V_i . This may be useful in cases where a local subsystem is known to be dissipative with respect to some w_i but no V_i has been computed.

3 Problem statement

As mentioned in Section 2, our aim is to certify the performance of the local subsystems in such a way that the

performance of the interconnected system is also certified. From now on, we will assume that all supply rates are quadratic. In particular,

$$\begin{aligned} w_0(d, e) &= \begin{bmatrix} d \\ e \end{bmatrix}^\top W \begin{bmatrix} d \\ e \end{bmatrix} \quad \text{is given, and} \\ w_i(u_i, y_i) &= \begin{bmatrix} u_i \\ y_i \end{bmatrix}^\top X_i \begin{bmatrix} u_i \\ y_i \end{bmatrix} \quad \text{for all } i \end{aligned} \quad (6)$$

where W and all X_i are real symmetric matrices. We are aiming to solve a feasibility problem of the form

$$\begin{aligned} \underset{X_{1:N}, V_{1:N}}{\text{minimize}} \quad & 0 \\ \text{subject to} \quad & (V_i, X_i) \in \mathcal{L}_i \quad \text{for } i = 1, \dots, N \\ & (X_1, \dots, X_N) \in \mathcal{G} \end{aligned} \quad (7)$$

Each \mathcal{L}_i constraint is *local* because it involves only the local supply rate X_i and storage function V_i . The \mathcal{G} constraint is *global* because it involves all the supply rates.

Before we define the \mathcal{L}_i and \mathcal{G} sets, first introduce the following conformal block partitions

$$W = \begin{bmatrix} W_{11} & W_{12} \\ W_{21} & W_{22} \end{bmatrix}, \quad X_i = \begin{bmatrix} X_i^{11} & X_i^{12} \\ X_i^{21} & X_i^{22} \end{bmatrix}$$

and the following block-diagonal matrices

$$X^{jk} = \begin{bmatrix} X_1^{jk} & & \\ & \ddots & \\ & & X_N^{jk} \end{bmatrix} \quad \text{for all } j, k \in \{1, 2\}$$

Recall that W is given while X_1, \dots, X_N are to be found. The local and global sets are defined as follows.

$$\begin{aligned} \mathcal{L}_i := \left\{ X_i, V_i \mid V_i \text{ is differentiable, } V_i(x_i) \geq 0, \right. \\ \left. \nabla V_i(x_i)^\top f_i(x_i, u_i) - \begin{bmatrix} u_i \\ y_i \end{bmatrix}^\top X_i \begin{bmatrix} u_i \\ y_i \end{bmatrix} \leq 0 \right. \\ \left. \text{for all } x_i, u_i \text{ and with } y_i = h_i(x_i, u_i) \right\} \end{aligned} \quad (8)$$

$$\mathcal{G} := \left\{ X_{1:N} \mid \sum_{i=1}^N H_i X_i H_i^\top - H_0 W H_0^\top \preceq 0 \right\} \quad (9)$$

Here, the constant matrices H_0, \dots, H_N are defined such that the following identity holds.

$$\begin{aligned} \sum_{i=1}^N H_i X_i H_i^\top - H_0 W H_0^\top = \\ \begin{bmatrix} M \\ I \end{bmatrix}^\top \begin{bmatrix} X^{11} & 0 & X^{12} & 0 \\ 0 & -W_{22} & 0 & -W_{21} \\ X^{21} & 0 & X^{22} & 0 \\ 0 & -W_{12} & 0 & -W_{11} \end{bmatrix} \begin{bmatrix} M \\ I \end{bmatrix} \end{aligned} \quad (10)$$

The local sets \mathcal{L}_i clearly certify the performance of the local subsystems. That is, each subsystem G_i is dissipative with respect to w_i if and only if $(X_i, V_i) \in \mathcal{L}_i$. The following result proves that any solution to (7) is a performance certificate for the interconnected system.

Proposition 1. *Consider the interconnected system described in Section 2, with supply rates are of the form (6). Further suppose that X_1, \dots, X_N and V_1, \dots, V_N are a solution to (7). Then the interconnected system is dissipative with respect to the supply rate w_0 .*

Proof. Multiplying the right-hand side of (10) on the left and right by $\begin{bmatrix} y^\top & d^\top \end{bmatrix}$ and $\begin{bmatrix} y^\top & d^\top \end{bmatrix}^\top$ respectively and making use of (3) and the definition (9), we obtain

$$\begin{bmatrix} u \\ e \\ y \\ d \end{bmatrix}^\top \begin{bmatrix} X^{11} & 0 & X^{12} & 0 \\ 0 & -W_{22} & 0 & -W_{21} \\ X^{21} & 0 & X^{22} & 0 \\ 0 & -W_{12} & 0 & -W_{11} \end{bmatrix} \begin{bmatrix} u \\ e \\ y \\ d \end{bmatrix} \leq 0 \quad (11)$$

Using this and the block diagonal structure of the X^{jk} , we can rewrite (11) as the following linear combination of quadratic forms.

$$\sum_{i=1}^N \begin{bmatrix} u_i \\ y_i \end{bmatrix}^\top X_i \begin{bmatrix} u_i \\ y_i \end{bmatrix} - \begin{bmatrix} d \\ e \end{bmatrix}^\top W \begin{bmatrix} d \\ e \end{bmatrix} \leq 0 \quad (12)$$

Adding to (12) each of the local dissipativity inequalities in (8), we obtain

$$\sum_{i=1}^N \nabla V_i(x_i)^\top f_i(x_i, u_i) - \begin{bmatrix} d \\ e \end{bmatrix}^\top W \begin{bmatrix} d \\ e \end{bmatrix} \leq 0$$

Hence, the interconnected system is dissipative with respect to the given supply rate w_0 . \blacksquare

The benefit of our formulation (7) is that verifying feasibility of a candidate point may be carried out in an efficient manner. The potentially complicated storage functions V_i only appear in the local constraints, so they can be checked separately and in parallel. Finally, membership of the \mathcal{G} set amounts to solving a global LMI. While this constraint does not decouple, it also does not depend on the storage functions.

In the next section, we will describe several different ways in which one may take advantage of the structure of (7) to decompose the optimization problem and thereby solve it efficiently.

4 Decomposition techniques

There are many ways to efficiently compute solutions to (7). The basic idea is that the constraints involving the V_i are independent, so one should take advantage of this fact in choosing a computational approach. As mentioned in Section 2, it is particularly important to take advantage of separability because the local storage functions V_i may need to be very complicated, and searching for a global storage function of adequate complexity may be computationally infeasible.

Projected Subgradient method. Dual decomposition and the projected subgradient method as presented in [7] was used in [18] to certify stability of an interconnected system and in [15] to verify the safety of a system using barrier certificates. Applying dual decomposition to (7), the projected subgradient method has the form:

1. Local optimization: given Λ^k , solve for each i :

$$(X_i^{k+1}, V_i^{k+1}) = \arg \min_{(X, V) \in \mathcal{L}_i} \text{trace}(\Lambda^k H_i X H_i^\top) + \rho \|X\|_F^2$$

where ρ is a regularization parameter.

2. If $(X_{1:N}^{k+1}, V_{1:N}^{k+1})$ is not feasible, update Λ^k as follows and return to step 1.

$$\Lambda^{k+1} = \mathbb{P} \left(\Lambda^k + \alpha_k \left(\sum_{i=1}^N H_i X_i^{k+1} H_i^\top - H_0 W H_0^\top \right) \right)$$

where \mathbb{P} projects onto the positive-semidefinite cone, and $\{\alpha_k\}$ is the stepsize sequence.

If the stepsize sequence $\{\alpha_k\}$ satisfies $\alpha_k > 0$, $\alpha_k \rightarrow 0^+$, and $\sum_k \alpha_k = \infty$, the subgradient method is guaranteed to converge to a solution, either after finitely many steps or in the limit. However, careful tuning of the stepsize schedule and regularization parameter is often required to achieve desirable performance [5].

Projection methods. A natural way of viewing (7) is that we are seeking to find a point that lies in the intersection of several sets. Note that \mathcal{G} is a convex set. If all the \mathcal{L}_i sets are convex as well, we may use projection methods. The simplest such method is the *alternating projection* method [9], which has the form:

1. For each i , project onto the local sets.

$$(X_i^{k+1/2}, V_i^{k+1}) = \arg \min_{(X, V) \in \mathcal{L}_i} \|X - X_i^k\|_F^2$$

2. project onto the global set.

$$X_{1:N}^{k+1} = \arg \min_{X_{1:N} \in \mathcal{G}} \sum_{i=1}^N \|X_i - X_i^{k+1/2}\|_F^2$$

3. If $(X_{1:N}^{k+1}, V_{1:N}^{k+1})$ is not feasible return to step 1.

We tested this method as well as *Dijkstra's method* [4], another popular projection method. These methods are guaranteed to converge monotonically, though the convergence rates may be very slow.

ADMM. The Alternating Direction Method of Multipliers (ADMM) is a kind of operator-splitting method, and is described in detail in [6]. This method has been shown to perform well in many practical engineering applications [6, 11, 12]. It can be used to solve problems of the form

$$\begin{aligned} & \text{minimize} && f(x) + g(z) \\ & \text{subject to} && Ax + Bz = c \end{aligned} \quad (13)$$

where x and z are vector decision variables. Our problem (7) may be put into this form by defining the following indicator functions:

$$\begin{aligned} \mathbb{I}_{\mathcal{L}_i}(X_i, V_i) &:= \begin{cases} 0 & (X_i, V_i) \in \mathcal{L}_i \\ \infty & \text{otherwise} \end{cases} \\ \mathbb{I}_{\mathcal{G}}(X_{1:N}) &:= \begin{cases} 0 & (X_1, \dots, X_N) \in \mathcal{G} \\ \infty & \text{otherwise} \end{cases} \end{aligned}$$

We may then write (7) as

$$\begin{aligned} & \text{minimize}_{X_{1:N}, Z_{1:N}, V_{1:N}} && \sum_{i=1}^N \mathbb{I}_{\mathcal{L}_i}(X_i, V_i) + \mathbb{I}_{\mathcal{G}}(Z_1, \dots, Z_N) \\ & \text{subject to} && X_i - Z_i = 0 \quad \text{for } i = 1, \dots, N \end{aligned}$$

and it is now in the canonical form (13). The f function and the constraints are separable, so the ADMM update takes on the following parallelized form [6].

1. X -updates: for each i , solve the local problem

$$(X_i^{k+1}, V_i^{k+1}) = \arg \min_{(X, V) \in \mathcal{L}_i} \|X - Z_i^k + U_i^k\|_F^2$$

2. Z -update: if $(X_{1:N}^{k+1}, V_{1:N}^{k+1})$ is not feasible, solve the global problem

$$Z_{1:N}^{k+1} = \arg \min_{Z_{1:N} \in \mathcal{G}} \left\| \sum_{i=1}^N (X_i^{k+1} - Z_i + U_i^k) \right\|_F^2$$

3. U -update: perform the following update and return to step 1.

$$U_i^{k+1} = X_i^{k+1} - Z_i^{k+1} + U_i^k$$

When applied to (13), ADMM is guaranteed to converge if f, g are closed, proper, and convex, and the Lagrangian has a saddle point [6].

5 Numerical simulations

In order to test and compare the different algorithms discussed in Section 4, we randomly generated $N = 50$ stable LTI subsystems, each with 3 states, 2 inputs, and 2 outputs. The subsystem state-space equations (1) then take the form

$$G_i : \begin{aligned} \dot{x}_i &= A_i x_i + B_i u_i \\ y_i &= C_i x_i + D_i u_i \end{aligned} \quad \text{for } i = 1, \dots, N$$

Each G_i was scaled so that its L_2 gain (equivalently, its H_∞ norm, or induced L_2 norm) was equal to 0.95. A random sparse interconnection matrix M was generated with 5% of its entries nonzero and scaled so that its spectral norm satisfies $\|M\| = 0.95$. The interconnected system has two inputs and two outputs. Due to the scalings mentioned above, the interconnected system will be stable and have L_2 gain less than 1. We then introduced further random scalings $\Psi = \text{diag}(\Psi_1, \dots, \Psi_N)$ and $\Phi = \text{diag}(\Phi_1, \dots, \Phi_N)$, and transformed the interconnection of Figure 1 into that of Figure 2. These final scalings do not change the closed-loop map from d to e , but they conceal the fact that the interconnection was constructed to satisfy a particular L_2 gain condition.




Figure 2: Scaling of the interconnected system of Figure 1 that leaves the closed-loop map unchanged.

As a proof-of-concept, we restricted our search to quadratic V_i functions. Therefore, the \mathcal{L}_i sets (8) are

$$\mathcal{L}_i := \left\{ X_i, P_i \mid P_i \succeq 0, \begin{bmatrix} A_i^\top P_i + P_i A_i & P_i B_i \\ B_i^\top P_i & 0 \end{bmatrix} - \begin{bmatrix} 0 & I \\ C_i & D_i \end{bmatrix}^\top X_i \begin{bmatrix} 0 & I \\ C_i & D_i \end{bmatrix} \preceq 0 \right\}$$

Certifying that the L_2 gain of the interconnected system is less than 1 amounts to using the supply rate

$$w_0(d, e) = \|d\|_2^2 - \|e\|_2^2 = \begin{bmatrix} d \\ e \end{bmatrix}^\top \begin{bmatrix} I & 0 \\ 0 & -I \end{bmatrix} \begin{bmatrix} d \\ e \end{bmatrix}$$

Figure 3 shows a typical convergence plot comparing the various methods discussed in Section 4. Each method was implemented in MATLAB using the CVX toolbox [8] to solve all the convex optimization problems. For each method, we plot the largest eigenvalue of the \mathcal{G} -constraint as a function of iteration count. When this value becomes negative, all eigenvalues are negative and the \mathcal{G} -constraint is satisfied; a feasible point has been found. The iterative methods were initialized using $\Lambda^0 = I$ and $X_i^0 = Z_i^0 = -I$, as applicable.

We now make a few remarks regarding these result. For the particular example plotted in Figure 3, the ADMM method converges in 15 iterations. The trace ends abruptly when the subsequent point is negative and feasibility has been attained.




Figure 3: Plot of the largest eigenvalue for five different iterative methods. Feasibility is achieved when all eigenvalues are negative (indicated by a terminal circle). ADMM converged in 15 iterations, while the other methods took longer or failed to converge after 1000 iterations.

Similar traces are obtained for the projected subgradient method, though the number of iterations required to find a feasible point turned out to be very sensitive to initial conditions, stepsize schedule $\{\alpha_k\}$, and regularization parameter ρ . With a stepsize of $\alpha_k = (100 + k)^{-1}$, a feasible point was found in 187 iterations, but when the stepsizes are increased slightly to $\alpha_k = (50 + k)^{-1}$, no feasible point was found after 1000 iterations. We used $\rho = 0.1$ for both cases.

Finally, Dykstra's method and the alternating projection method both exhibited a monotonically decreasing behavior typical of projection methods; they steadily approached the boundary of the feasible set without ever penetrating it.

Note that in Figure 3, the x -axis measures iterations rather than time or floating-point operations. Iteration count is a fair metric in this case, because all methods compared have a similar structure. Specifically, each iteration consists of parallelizable local steps involving the \mathcal{L}_i constraints, and a global step involving the \mathcal{G} constraint. There is a minor difference in the global step, however. The subgradient method computes a global projection onto the positive-semidefinite cone, while the other methods compute a projection onto \mathcal{G} , which is described by a global LMI. One could argue that the latter requires more computational effort than the former.

In order to test the robustness of ADMM, we generated 1000 random instances of the interconnected system described earlier in this section. All parameters including the interconnection topology were randomized. ADMM was applied to each system, and the number of iterations required to certify L_2 performance were recorded. Figure 4 shows a cumulative frequency plot of the result. Only one of the 1000 systems tested required more than 40 iterations (it required 74 iterations), and 90% of systems tested required 16 iterations or fewer.

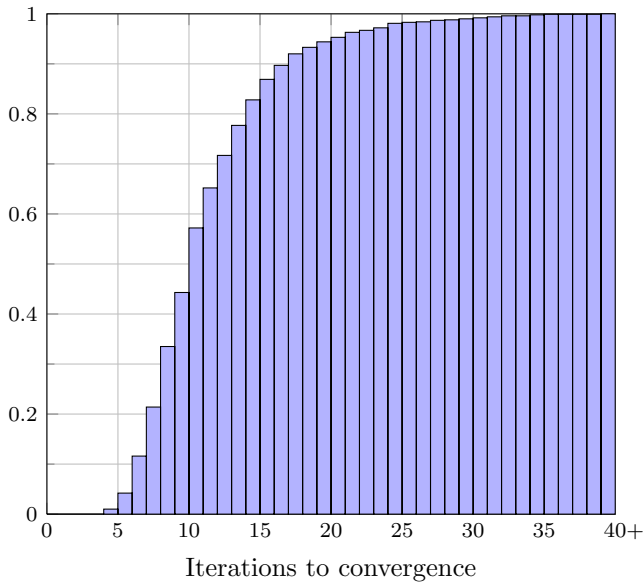


Figure 4: Cumulative plot showing the fraction of 1000 total trials that required at most a given number of iterations to find a feasible point using ADMM. For example, the fastest trials found a feasible point in 4 iterations. Also, 90% of trials succeeded in 16 iterations or fewer.

These results suggest that ADMM, even in its most basic implementation, is a fast and reliable way of solving large distributed performance certification problems, both in an absolute sense, and as compared to sub-gradient or projection methods.

6 Extensions and future directions

Several extensions are possible, and are likely necessary to render the method useful for its intended purpose of certifying performance of large interconnections of nonlinear systems. We discuss two possibilities below.

Polynomials and SOS. In Section 5, we assumed the V_i were quadratic. This may not be adequate for certifying the dissipativity of a general nonlinear system, as in equation (4)–(5). A sufficient condition, suitable for the case when f, h are polynomials and V is drawn from a finite-dimensional subspace of polynomial functions, is to require that the expressions $V(x)$ and

$w(u, h(x, u)) - \nabla V^\top f(x, u)$ are sum-of-squares in the variables x and (x, u) respectively. Deciding if such a V exists can be determined by a semidefinite program.

Extension to IQCs. Rather than requiring dissipativity with respect to a quadratic supply rate, we may want to allow for integral quadratic constraints (IQCs). This allows for very general frequency-dependent performance measures, including robustness analysis [10]. As in the dissipativity framework, the conventional approach of fixing the local storage functions was adapted to IQCs in [16]. We now outline how the approach of Sections 3 and 4 may also be extended to IQCs.

If $(\bar{A}, \bar{B}, \bar{C}, \bar{D})$ is the realization of a stable LTI system Ψ , then the existence of a positive semidefinite $V(x, \eta)$ such that

$$\begin{aligned} (\nabla_x V)^\top f(x, u) + (\nabla_\eta V)^\top \begin{pmatrix} \bar{A}\eta + \bar{B} \begin{bmatrix} u \\ y \end{bmatrix} \\ \bar{C}\eta + \bar{D} \begin{bmatrix} u \\ y \end{bmatrix} \end{pmatrix} \\ \leq \begin{pmatrix} \bar{C}\eta + \bar{D} \begin{bmatrix} u \\ y \end{bmatrix} \end{pmatrix}^\top X \begin{pmatrix} \bar{C}\eta + \bar{D} \begin{bmatrix} u \\ y \end{bmatrix} \end{pmatrix} \end{aligned} \quad (14)$$

for all x, u , and $y = h(x, u)$ ensures that the dynamical system $\dot{x} = f(x, u)$, $y = h(x, u)$ satisfies the IQC defined by $\Pi = \Psi^* X \Psi$. We recover traditional dissipativity in the special case where Ψ is static.

Using such IQCs for each subsystem, the global certification condition (9) takes on the form

$$\begin{bmatrix} M \\ I \end{bmatrix}^\top \begin{bmatrix} \Pi^{11} & 0 & \Pi^{12} & 0 \\ 0 & -W_{22} & 0 & -W_{21} \\ \Pi^{21} & 0 & \Pi^{22} & 0 \\ 0 & -W_{12} & 0 & -W_{11} \end{bmatrix} \begin{bmatrix} M \\ I \end{bmatrix} \preceq -\varepsilon I$$

where the Π^{jk} and W_{jk} are functions of frequency ω , and the inequality must hold for all $\omega \in \mathbb{R}$. Such a constraint can be made tractable by sampling at finitely many points, or via the KYP lemma.

While these extensions affect the sets (4)–(5), the optimization problem (7) has the same structural form. Hence, the ADMM algorithm as described in Section 4 can be applied, and is the subject of future work.

7 Acknowledgements

The authors would like to thank Murat Arcak for helpful discussions. The second author would also like to thank Matt Kranning, Sanjay Lall, and Stephen Boyd.

This work was supported by NASA under Grant No. NRA NNX12AM55A entitled ‘‘Analytical Validation Tools for Safety Critical Systems Under Loss-of-Control Conditions’’, Dr. Christine Belcastro technical monitor. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of NASA.

References

- [1] J. Anderson and A. Papachristodoulou. A network decomposition approach for efficient sum of squares programming analysis. In *American Control Conference*, pages 4492–4497, 2010.
- [2] J. Anderson, A. Teixeira, H. Sandberg, and A. Papachristodoulou. Dynamical system decomposition using dissipation inequalities. In *IEEE Conference on Decision and Control*, pages 211–216, 2011.
- [3] M. Arcak and E. D. Sontag. A passivity-based stability criterion for a class of biochemical reaction networks. *Mathematical Biosciences and Engineering*, 5(1):1–19, 2008.
- [4] H. Bauschke and J. Borwein. Dykstra’s alternating projection algorithm for two sets. *Journal of Approximation Theory*, 79(3):418–443, 1994.
- [5] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty. *Nonlinear Programming: Theory and Algorithms*. Wiley-InterScience, 3rd edition, 2006.
- [6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [7] S. Boyd, L. Xiao, A. Mutapcic, and J. Mattingley. Notes on decomposition methods, Feb. 2007. http://www.stanford.edu/class/ee364b/notes/decomposition_notes.pdf.
- [8] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.0 beta. <http://cvxr.com/cvx>, Sept. 2013.
- [9] L. G. Gubin, B. T. Polyak, and E. V. Raik. The method of projections for finding the common point of convex sets. *USSR Computational Mathematics and Mathematical Physics*, 7(6):1–24, 1967.
- [10] A. Megretski and A. Rantzer. System analysis via integral quadratic constraints. *IEEE Transactions on Automatic Control*, 42(6):819–830, 1997.
- [11] J. F. C. Mota, J. M. F. Xavier, P. M. Q. Aguiar, and M. Püschel. Distributed ADMM for model predictive control and congestion control. In *IEEE Conference on Decision and Control*, pages 5110–5115, 2012.
- [12] M. K. Ng, P. Weiss, and X. Yuan. Solving constrained total-variation image restoration and reconstruction problems via alternating direction methods. *SIAM Journal on Scientific Computing*, 32(5):2710–2736, 2010.
- [13] H. Sandberg. Hankel-norm-based lumping of interconnected linear systems. In *Proceedings of MATHMOD*, 2009.
- [14] N. Sandell, Jr., P. Varaiya, M. Athans, and M. Saffonov. Survey of decentralized control methods for large scale systems. *IEEE Transactions on Automatic Control*, 23(2):108–128, 1978.
- [15] C. Sloth, G. J. Pappas, and R. Wisniewski. Compositional safety analysis using barrier certificates. In *ACM international conference on Hybrid Systems: Computation and Control*, pages 15–24, 2012.
- [16] E. Summers. *Performance analysis of nonlinear systems combining integral quadratic constraints and sum-of-squares techniques*. PhD thesis, University of California, Berkeley, 2012.
- [17] E. Summers, A. Chakraborty, W. Tan, U. Topcu, P. Seiler, G. Balas, and A. Packard. Quantitative local L_2 -gain and reachability for nonlinear systems. *International Journal of Robust and Nonlinear Control*, 23(10):1115–1135, 2013.
- [18] U. Topcu, A. Packard, and R. M. Murray. Compositional stability analysis based on dual decomposition. In *IEEE Conference on Decision and Control*, pages 1175–1180, 2009.
- [19] U. Topcu, A. Packard, and P. Seiler. Local stability analysis using simulations and sum-of-squares programming. *Automatica*, 44(10):2669–2675, 2008.
- [20] J. C. Willems. Dissipative dynamical systems part I: General theory. *Archive for Rational Mechanics and Analysis*, 45(5):321–351, 1972.
- [21] J. C. Willems. Dissipative dynamical systems part II: Linear systems with quadratic supply rates. *Archive for Rational Mechanics and Analysis*, 45(5):352–393, 1972.