

# Resistance Distance and Control Performance for bittide Synchronization

Sanjay Lall<sup>1</sup>

Călin Cașcaval<sup>2</sup>

Martin Izzard<sup>2</sup>

Tammo Spalink<sup>2</sup>

## Abstract

We discuss control of bittide distributed systems, which are designed to provide logical synchronization between networked machines by observing data flow rates between adjacent systems at the physical network layer and controlling local reference clock frequencies. We analyze the performance of approximate proportional-integral control of the synchronization mechanism and develop a simple continuous-time model to show the resulting dynamics are stable for any positive choice of gains. We then construct explicit formulae to show that closed-loop performance measured using the  $L_2$  norm is a product of two terms, one depending only on resistance distances in the graph, and the other depending only on controller gains.

## 1 Introduction

In this paper, we discuss control of the bittide synchronization mechanism for distributed computing. The purpose of this mechanism is to provide all of the machines on a network with shared logical time. This notion of time does not have to match physical time. Instead, the discrete clocks of the machines on the network are tied together in *logical synchronization*. This is distinct from *physical synchronization*, where the processor clocks are kept synchronized to physical time [4, 20]. Application processes on this system coordinate their actions using logical time, and do not need to reference physical time. The bittide control system maintains perfect logical synchronization using imperfect physical synchronization.

We view the network as an undirected graph where each edge represents a pair of data links between nodes (machines), one link in each direction. The synchronization mechanism operates at the physical layer of the network as follows. Frames of data received from each link are appended at the tail of per-link queues called *elastic buffers*. One frame is removed from the head of each elastic buffer every local clock cycle and consumed by the

corresponding compute core. One frame is also transmitted on every outgoing link during every cycle. This alignment of receive and transmit at each node reveals the relative frequency between neighboring nodes. If the elastic buffer at a node starts to fill up, then it must have a lower clock frequency than its neighbor and vice versa.

For a network with  $n$  nodes, if node  $i$  has  $d_i$  neighbors, then it has  $d_i$  elastic buffers, one per node. Frames are removed from all of the  $d_i$  elastic buffers simultaneously, driven by the local clock. At each node  $i$ , the local clock is driven by a physical oscillator, with uncorrected frequency  $\omega_i^u$ . The uncorrected frequencies at nodes will differ slightly in practice, and so additional correction is necessary to ensure system stability. Each node includes a feedback control system which measures the occupancy of all local elastic buffers and adds a correction  $c_i$  to the local oscillator frequency such that it oscillates at frequency  $\omega_i = c_i + \omega_i^u$ . The purpose of the control system is to prevent the elastic buffers from overflowing or underflowing, even though the  $\omega_i^u$  are not known exactly. Full details of this mechanism are presented in [18], where a mathematical model called the *abstract frame model* (AFM) is developed.

Our focus in this paper is the use of proportional and proportional-integral control for bittide synchronization. To that end we approximate the abstract frame model with a simple linear model, removing the effects of sampling and quantization. We present simulations illustrating this approximation, and analyze the mathematical properties of the resulting linear system. We show that it is stable, and that certain closed-loop performance metrics can be expressed in terms of the resistance matrix of the graph. Performance is measured using the  $L_2$  norm of the buffer occupancies and frequency deviations, for which we give exact formulae, in terms of the resistance distances in the graph and the controller gains. These results directly relate the connectivity of the graph to the performance properties of the bittide system.

**Prior work.** The synchronization mechanism of bittide was first proposed in [24]. The abstract frame model for the system was developed in [18], where a detailed description of the dynamic behavior of the system was given. Another widely-used synchronous network mechanism is SONET [27]. The use of coupled-oscillators to model synchronization originates with Winfree [29].

<sup>1</sup>S. Lall is with the Department of Electrical Engineering at Stanford University, Stanford, CA 94305, USA, and is a Visiting Researcher at Google. [lall@stanford.edu](mailto:lall@stanford.edu)

<sup>2</sup>Călin Cașcaval, Martin Izzard, and Tammo Spalink are with Google.

In this paper, we approximate the bittide mechanism with a linear model. Our focus is on proportional-integral (PI) control, but the corresponding model with purely proportional control is the widely studied Laplacian dynamics, which has been extensively studied in the literature, with applications including models of flocking [14, 23], Markov chain averaging models [2, 13], congestion control protocols [15], power networks [5, 25], vehicle platooning [26], and consensus [22]. Nonlinear versions of Laplacian dynamics have been studied in [28], and papers addressing PI control of Laplacian dynamics include [1, 3, 8].

## 2 Modeling

We consider an oriented graph with  $n$  nodes and  $m$  edges. Although the graph is undirected, each edge has an *orientation* used purely to define the sign convention. For an oriented graph we define the incidence matrix  $B \in \mathbb{R}^{n \times m}$  by

$$B_{il} = \begin{cases} 1 & \text{if node } i \text{ is the source of edge } l \\ -1 & \text{if node } i \text{ is the target of edge } l \\ 0 & \text{otherwise} \end{cases}$$

and this defines a numbering of the edges  $l = 1, \dots, k$ . The matrix  $L = BB^\top$  is the *Laplacian* matrix of the graph. We assume the graph is connected. It is then a standard result that  $L$  has rank  $n - 1$ . There is exactly one zero eigenvalue, with corresponding eigenvector  $\mathbf{1}$ . See for example [10]. Choose  $U_1$  to complete the basis, so that  $U \in \mathbb{R}^{n \times n}$  is an orthogonal matrix such that

$$U = [U_1 \quad \mathbf{1}/\sqrt{n}]$$

Then we can write  $L$  in these coordinates so that

$$U^\top L U = \begin{bmatrix} \hat{L} & 0 \\ 0 & 0 \end{bmatrix}$$

where  $\hat{L} \in \mathbb{R}^{(n-1) \times (n-1)}$  is positive definite.

**The abstract frame model.** We briefly summarize the abstract frame model [18]. We have  $n$  nodes. At each node  $i$  there is a clock, whose value  $\theta_i \in \mathbb{R}$  is called the *clock phase*. We say that  $\theta_i$  measures local time at node  $i$ , in units called *local ticks*. The rate of change for  $\theta_i$  is called the *frequency* of node  $i$ , denoted by  $\omega_i$ . Every time  $t$  at which the phase  $\theta_i(t)$  is an integer, node  $i$  sends a data frame to each of its neighbors. The number of frames in the elastic buffer at node  $j$  associated with the link from node  $i$  at time  $t$  is called the buffer *occupancy*, denoted by  $\beta_{ij}(t)$ . One can show that

$$\beta_{ij}(t) = \lfloor \theta_i(t - l_{ij}) \rfloor - \lfloor \theta_j(t) \rfloor + \lambda_{ij}$$

where  $\lambda_{ij} \in \mathbb{Z}$  is a constant, and  $l_{ij}$  is the latency of the link from  $i$  to  $j$ . Every  $p$  local ticks, the controller at

node  $i$  measures the buffer occupancies  $\beta_{ji}$  at that node. After a delay of  $d$  local ticks, the controller sets the frequency correction  $c_i$ . The delay parameter  $d$  specifies the time required by the controller to process the measurements and choose the frequency correction and includes the time for the frequency change to take effect on the physical oscillator. The dynamic model is as follows. For all  $t \geq 0$ ,  $i \in \mathcal{V}$ , and  $k \in \mathbb{Z}_+$ ,

$$\begin{aligned} \dot{\theta}_i(t) &= c_i^k + \omega_i^u \quad \text{for } t \in [s_i^k, s_i^{k+1}) \\ \beta_{ji}(t) &= \lfloor \theta_j(t - l_{ji}) \rfloor - \lfloor \theta_i(t) \rfloor + \lambda_{ji} \\ \theta_i(t_i^k) &= \theta_i^0 + kp \\ \theta_i(s_i^k) &= \theta_i^0 + kp + d \\ y_i^k &= \{(j, \beta_{ji}(t_i^k)) \mid j \in \text{neighbors}(i)\} \\ \xi_i^{k+1} &= f_i^{\text{disc}}(\xi_i^k, y_i^k) \\ c_i^k &= g_i^{\text{disc}}(\xi_i^k, y_i^k) \end{aligned} \tag{1}$$

Because  $\theta_i$  is an increasing function, the third and fourth equations above uniquely determine the sampling times  $t_i^k$  and the hold times  $s_i^k$ . The controller is given by a discrete-time state-space system  $f_i^{\text{disc}}, g_i^{\text{disc}}$  at each node which maps the history of these measurements to the correction  $c_i$ . The initial conditions of the model are

$$\theta_i(t) = \begin{cases} \theta_i^0 + \omega_i^{(-2)}t & \text{for } t \in [t^e, 0] \\ \theta_i^0 + \omega_i^{(-1)}t & \text{for } t \in [0, d/\omega_i^{(-1)}] \end{cases} \tag{2}$$

These conditions are determined by initial frequencies  $\omega_i^{(-1)} > \omega^{\min}$  and  $\omega_i^{(-2)} > \omega^{\min}$ , and initial clock phases  $\theta_i^0 \in \mathbb{R} + \mathbb{Z}$ . We are also given the initial buffer occupancies  $\beta_{ji}^0 \in \mathbb{Z}_+$ , which together with (2) determine the constants  $\lambda_{ji}$  such that

$$\beta_{ji}(t) - \beta_{ji}^0 = \lfloor \theta_j(t - l_{ji}) \rfloor - \lfloor \theta_j(-l_{ji}) \rfloor - (\lfloor \theta_i(t) \rfloor - \lfloor \theta_i(0) \rfloor)$$

A controller is called *admissible* if

$$g_i^{\text{disc}}(\xi, y) + \omega_i^u > \omega^{\min}$$

for all  $i, k$  and all measurements  $y$  and controller states  $\xi$ . This ensures that  $\omega_i(t) > \omega^{\min}$  for all  $i, t$ . The time  $t^e < 0$  is called the *epoch*, and it must satisfy

$$t^e \leq -(l_{ji} + d/\omega^{\min}) \text{ for all } i, j \in \mathcal{V}$$

We have shown in [18] that the abstract frame model has a unique solution under these conditions.

### 2.1 An approximate model

Our goal is to design a controller using a model for the system that is as close to the AFM as possible. However, in order to mathematically analyze and validate the controller, we need to use a simplified model. We perform two important simplifications. Section 4 includes simulations of both the AFM and the simplified model for comparison.

**Continuous-time approximation.** The first simplification is that we design a continuous-time controller, instead of designing a discrete-time controller that uses sampled-data. Continuous-time control is not practically possible for bittide, and therefore in implementations we need to discretize the controller and subsequently validate that this does not adversely affect performance. With a sufficiently fast sampling rate, standard approaches prove and demonstrate that this is an effective design methodology.

We will use a continuous-time controller of the form

$$\begin{aligned} \frac{d\xi_i}{d\theta_i} &= f_i(\xi_i, y_i) \\ c_i(t) &= g_i(\xi_i, y_i) \end{aligned} \quad (3)$$

defined by functions  $f_i$  and  $g_i$  at node  $i$ . Notice here that the independent variable is  $\theta_i$ , not  $t$ . This captures the dependence of the controller dynamics on the local oscillator frequency, which arises because the clock that drives the discrete-time controller is provided by the oscillator at the node. We write this in terms of  $t$  as follows. Since  $\dot{\theta}_i = \omega_i$ , we have

$$\begin{aligned} \dot{\xi}_i(t) &= \omega_i(t) f_i(\xi_i(t), y_i(t)) \\ c_i(t) &= g_i(\xi_i(t), y_i(t)) \end{aligned} \quad (4)$$

The nodes in bittide cannot exactly execute arbitrary dynamic control since they do not have access to a perfect physical time reference. Any integration or differentiation performed by the controller will be scaled by the current clock frequency  $\omega_i$ , which the controller cannot measure. Since the clock frequency is determined by the controller, this introduces a nonlinear feedback into the system. The magnitude of this effect depends on the range of frequency variation experienced at the node. This is in practice determined by the accuracy of the physical oscillators used. If the oscillator frequency is accurate to within a relative error of  $\alpha$ , then only correspondingly small relative corrections are required by the controller, and therefore the term  $\omega_i$  in equation (4) will be constant to within a relative error  $\alpha$  also. In a practical bittide implementation we might see  $\alpha < 10^{-5}$ , which is substantially below the gain margin typically used in control design. Therefore this much uncertainty in the controller parameters may be safely ignored. Hence we can approximate the controller dynamics by

$$\begin{aligned} \dot{\xi}_i &= \omega^c f_i(\xi_i, y_i) \\ c_i(t) &= g_i(\xi_i, y_i) \end{aligned} \quad (5)$$

where  $\omega^c$  is a constant approximation to the frequency (and hence not node specific.)

**Quantization.** The second simplification that we apply to the model is that we remove the quantization of frames. That is, instead of enforcing the physical property that the elastic buffer contains an integer number of

frames, we modify the model to allow the buffer occupancy to be non-integral. Replacing discrete-frames by a continuum results in a so-called *fluid model*, often used in analysis of stochastic models of queuing systems [17]. This corresponds to replacing the expression for occupancy  $\beta_{ji}$  by

$$\beta_{ji}(t) = (\lfloor a\theta_j(t - l_{ji}) \rfloor - \lfloor a\theta_i(t) \rfloor) / a + \lambda_{ji}$$

and taking the limit as  $a \rightarrow \infty$ . Subject to mild technical conditions, if the controller is linear, then solution trajectories converge in  $L_\infty$ . The limiting dynamics has an approximate occupancy given by

$$\beta_{ji}(t) = \theta_j(t - l_{ji}) - \theta_i(t) + \lambda_{ji}$$

We can expect this to be a good approximation if the frames are moving sufficiently quickly through the system in comparison to the timescale of the controller.

**Zero delays.** A further approximation that we make in this paper is that the computation delays  $d$  and the latencies  $l_{ji}$  are small enough to be neglected, and set to zero. In practice this may or may not be the case; some links have very long latencies. In other cases, such as between machines in a datacenter, the latencies are very short.

The assumptions that quantization, loop delays, and sampling may be ignored for the purposes of design, are well-studied in the literature and frequently used in practice. Mathematical techniques for handling the error due to quantization exist, for example [9]. The problem of multi-rate sampled-data systems has also been studied, for example [12, 19]. Systems with multiple delays are analyzed in [21]. In this case we have all of these phenomena together with state dependent sampling rates and so further analytical developments are required. Thus, in this paper we do not validate the assumption that the error induced by these approximations is small. We have as yet performed neither the exhaustive numerical simulations nor the mathematical analysis required to do so. However, we present some simple numerical simulations below that indicate the approximation error is very small in some cases of interest.

Combining these approximations and applying them to the abstract frame model gives the following model:

$$\begin{aligned} \dot{\theta}_i(t) &= c_i(t) + \omega_i^u \\ \beta_{ji}(t) - \beta_{ji}^0 &= \theta_j(t) - \theta_j^0 - (\theta_i(t) - \theta_i^0) \\ y_i(t) &= \{(j, \beta_{ji}(t)) \mid j \in \text{neighbors}(i)\} \\ \dot{\xi}_i(t) &= \omega^c f_i(\xi_i(t), y_i(t)) \\ c_i(t) &= g_i(\xi_i(t), y_i(t)) \end{aligned} \quad (6)$$

### 3 Controllers

The objective of the control design is to ensure that the elastic buffers neither overflow nor underflow. That is,

we must ensure that  $0 \leq \beta_{ji}(t) \leq \beta^{\max}$  for all  $t \geq 0$  and  $i, j \in \mathcal{V}$ . The controller determines the frequency correction  $c_i(t)$ , and must ensure that  $\omega^{\min} < c_i(t) + \omega_i^u < \omega^{\max}$ , where  $\omega^{\min} > 0$  and  $\omega^{\max}$  are the physical limits of the oscillator.

To do this, we initialize the buffer occupancies in the middle of the buffer, so that  $\beta_{ji}^0 = \beta^{\max}/2$  ( $\beta^{\max}$  is even by construction.) We will then construct a feedback signal using the offset

$$\bar{\beta}_{ji}(t) = \beta_{ji}(t) - \beta_{ji}^0$$

For convenience, we also define  $\bar{\theta}_i(t) = \theta_i(t) - \theta_i^0$ . We will consider proportional and proportional-integral controllers. Since each node may have a different number of neighbors, we must aggregate the measurements of elastic buffer occupancy at each node, and we do this simply by summing the occupancies, so that the controller at node  $i$  uses the feedback signal

$$r_i = \sum_{j \in \text{neighbors}(i)} \bar{\beta}_{ji}(t)$$

Other choices for the measurement are possible. For example, one might use the average buffer occupancy at a node, or the maximum occupancy at a node. At this point in the analysis it seems as though any of these choices might work well. However it turns out that using the sum of occupancies results in the dynamics of the closed-loop system being approximately equal to the well-known Laplacian dynamics on the graph, which has many good properties, including guaranteed stability with proportional control.

So far we have used notation  $\beta \in \mathbb{R}^{n \times n}$  to denote elastic buffer occupancy, with  $\beta_{ji}$  the occupancy of the buffer at node  $i$  associated with edge  $(i, j)$ , and  $\beta_{ij}$  the occupancy of the buffer at the other end of the edge. Under the zero-latency assumption, we have  $\bar{\beta}_{ji} = -\bar{\beta}_{ij}$ . We therefore define the *relative buffer occupancy*  $\delta \in \mathbb{R}^m$  by

$$\delta_l = \bar{\beta}_{ji} \text{ if } i \text{ is the source of edge } l$$

and can replace  $\bar{\beta} \in \mathbb{R}^{n \times n}$  by the equivalent representation  $\delta \in \mathbb{R}^m$ . We can now write the above dynamics as follows.

$$\begin{aligned} \dot{\bar{\theta}}(t) &= c(t) + \omega^u \\ \delta(t) &= -B^T \bar{\theta}(t) \\ r(t) &= B\delta(t) \\ \dot{\xi}_i(t) &= \omega^c f_i(\xi_i(t), r_i(t)) \\ c_i(t) &= g_i(\xi_i(t), r_i(t)) \end{aligned} \quad (7)$$

## 4 Proportional-integral control

We approximate a proportional-integral controller as follows. The controller is

$$\begin{aligned} \dot{\xi}(t) &= \omega^c r(t) & \xi(0) &= 0 \\ c(t) &= k_I \xi(t) + k_P r(t) \end{aligned}$$

Choosing state coordinates  $x_1 = \bar{\theta}$  and  $x_2 = \xi/\omega^c$ , we have, with  $x = (x_1, x_2)$ ,

$$\begin{aligned} \dot{x} &= Ax + B_2 \omega^u & x(0) &= 0 \\ \omega &= C_1 x + D_1 \omega^u \\ \delta &= C_2 x \end{aligned} \quad (8)$$

where for convenience  $a = k_P$ ,  $b = \omega^c k_I$  and

$$\begin{aligned} A &= \begin{bmatrix} -aL & bI \\ -L & 0 \end{bmatrix} & B_2 &= \begin{bmatrix} I \\ 0 \end{bmatrix} \\ C_1 &= \begin{bmatrix} -aL & bI \end{bmatrix} & D_1 &= I \\ C_2 &= \begin{bmatrix} -B^T & 0 \end{bmatrix} \end{aligned}$$

Figure 1 shows a comparison of the abstract frame model (1) and the continuous-time ordinary differential equation model (8). This is for a graph with 3 nodes, with bidirectional links connecting every pair of nodes. The system parameters are  $k_P = 3 \times 10^{-5}$ ,  $k_I = 2 \times 10^{-9}$ ,  $l_{ji} = 500$ ,  $p = 1000$ ,  $d = 100$ ,  $\theta_i^0 = 0.1$ ,  $\omega^{(-2)} = \omega^u$ ,  $\omega^{(-1)} = \omega^u$ , and  $\omega^c = 1$ . This simulation shows that, for this choice of parameters, the continuous-time model approximates the behavior of the AFM, at least qualitatively. A more in-depth analysis of this approximation shows that the effects of the quantization become insignificant for large buffer occupancies. As might be expected, the effects of the latencies and computational delays are also negligible if the control gains are not too large. We will not pursue this comparison further here, but assume for the purposes of this paper that the parameters of the system are such that the AFM behavior is well-approximated by the differential equation model.

## 5 System Behavior

The proofs of the results in this section may be found in the Appendix. Using the change of coordinates

$$x = \begin{bmatrix} U_1 & 0 & U_2 & 0 \\ 0 & U_1 & 0 & U_2 \end{bmatrix} \hat{x}$$

we have the dynamics

$$\begin{aligned} \dot{\hat{x}} &= \begin{bmatrix} -a\hat{L} & bI & 0 & 0 \\ -\hat{L} & 0 & 0 & 0 \\ 0 & 0 & 0 & bI \\ 0 & 0 & 0 & 0 \end{bmatrix} \hat{x} + \begin{bmatrix} U_1^T \\ 0 \\ U_2^T \\ 0 \end{bmatrix} \omega^u \\ \omega &= [-aLU_1 \quad bU_1 \quad 0 \quad bU_2] \hat{x} + \omega^u \\ \delta &= [-B^T U_1 \quad 0 \quad 0 \quad 0] \hat{x} \end{aligned}$$

In these coordinates, we have immediately that  $\hat{x}_4 = 0$ , since the initial conditions are  $\hat{x} = 0$ . This then implies that

$$\hat{x}_3(t) = n^{\frac{1}{2}} \omega^{\text{avg}} t$$

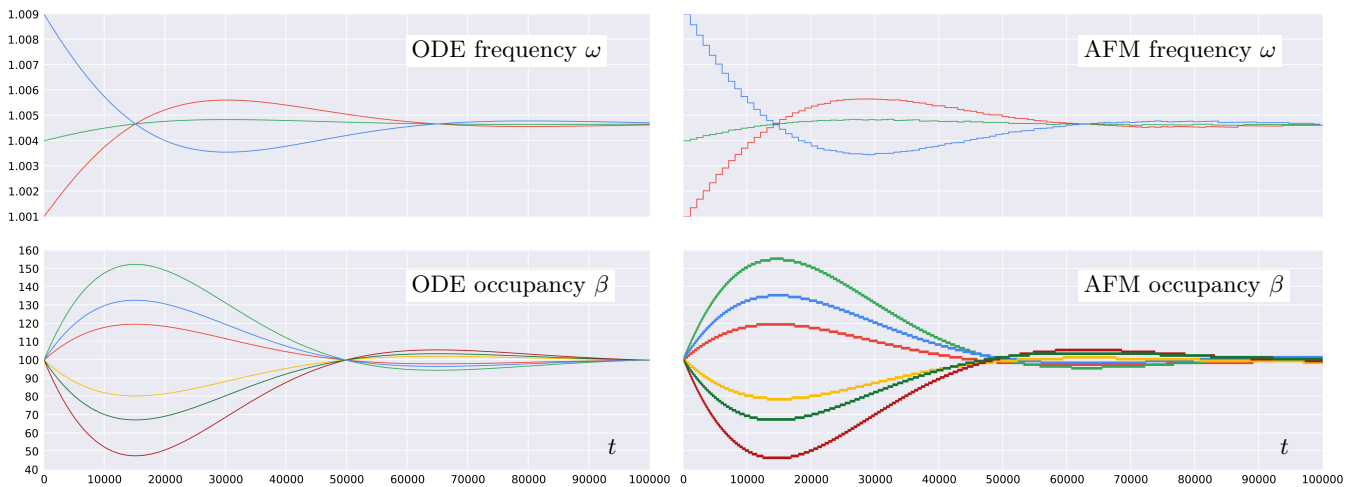


Figure 1: Comparison of trajectories from two different models

We now have the remaining dynamics

$$\begin{aligned}\dot{\tilde{x}} &= \hat{A}\tilde{x} + \hat{B}_2\omega^u \\ \omega &= \hat{C}_1\tilde{x} + \hat{D}_1\omega^u \\ \delta &= \hat{C}_2\tilde{x}\end{aligned}$$

where  $\tilde{x} = (\hat{x}_1, \hat{x}_2)$  and

$$\begin{aligned}\hat{A} &= \begin{bmatrix} -a\hat{L} & bI \\ -\hat{L} & 0 \end{bmatrix} & \hat{B}_2 &= \begin{bmatrix} U_1^\top \\ 0 \end{bmatrix} \\ \hat{C}_1 &= [-aLU_1 & bU_1] & \hat{D}_1 &= I \\ \hat{C}_2 &= [-B^\top U_1 & 0] & \hat{C} &= \begin{bmatrix} \hat{C}_1 \\ \hat{C}_2 \end{bmatrix}\end{aligned}\quad (9)$$

We would like to show that these dynamics are stable, which we state formally here.

**Theorem 1.** *Suppose  $a > 0$  and  $b > 0$ . Then the matrix  $\hat{A}$  in (9) is Hurwitz.*

Since  $\hat{A}$  is stable,  $\tilde{x}$  converges to a steady-state value  $\tilde{x}(t) \rightarrow \tilde{x}^{ss}$  as  $t \rightarrow \infty$ , and this is

$$\tilde{x}^{ss} = -\hat{A}^{-1}\hat{B}_2\omega^u = \begin{bmatrix} 0 \\ \frac{1}{b}U_1^\top \end{bmatrix} \omega^u \quad (10)$$

We can now make the following observations regarding the behavior of this system, which follow immediately from the above representation.

- i) The phase  $\bar{\theta}$  is the sum of a transient term which tends to zero and a linearly growing term whose growth rate is determined by the average frequency error. We have

$$\begin{aligned}\bar{\theta}(t) &= U_1\hat{x}_1 + U_2\hat{x}_3 \\ &= U_1\tilde{x}_1 + \omega^{avg}t\mathbf{1}\end{aligned}$$

where  $\tilde{x}_1 \rightarrow 0$  as  $t \rightarrow \infty$ .

- ii) The sum of the phases grows linearly with time, since  $\mathbf{1}^\top U_1 = 0$ .

- iii) The frequency  $\omega = \dot{\bar{\theta}}$  has an invariant sum

$$\sum_{i=1}^n \omega_i(t) = n\omega^{avg} \quad \text{for all } t$$

and all frequencies converge to the average frequency, that is  $\omega_i(t) \rightarrow \omega^{avg}$  as  $t \rightarrow \infty$ .

- iv) The relative buffer occupancies tend to zero, that is  $\delta(t) \rightarrow 0$  as  $t \rightarrow \infty$ .

## 5.1 Performance and Resistance Distance

We now turn to performance measures of this controller. Specifically, one of the primary controller objectives is to keep  $\delta$  small. This means that the buffer occupancies will remain close to the middle of the buffer  $\beta^0$ , reducing the chance of the buffers overflowing or underflowing. We consider here the 2-norm as a measure of the magnitude of  $\delta$ . It will turn out that this quantity is related to the connectivity graph of the system, and so this affords a design strategy for the network topology. We can choose topologies such that the norm of the buffer occupancy is small. However, we note that this is simply a heuristic for the specific objective of preventing buffer overflows and underflows.

For convenience let  $\omega^{ss}$  be the steady-state frequency, given by  $\omega^{ss} = \omega^{avg}\mathbf{1}$ . We refer to the quantity  $\omega - \omega^{avg}\mathbf{1}$  as the *frequency deviation*, and the first result concerns its norm.

**Theorem 2.** *For the dynamics as above, we have*

$$\|\omega - \omega^{ss}\|^2 = \frac{\omega^u{}^\top L^\dagger \omega^u}{2a}$$

The second result concerns the norm of the buffer occupancies.

**Theorem 3.** For the dynamics as above, we have

$$\|\delta\|^2 = \frac{\omega^u \top L^\dagger \omega^u}{2ab}$$

**The effect of the controller.** The above theorems separate the effects on performance of the controller from the effects of the graph. We see that increasing the proportional gain  $a = k_P$  improves both frequency and occupancy performance. However, increasing the (scaled) integral gain  $b = \omega^c k_I$  improves occupancy performance, but does not change the norm of the frequency deviation.

**The effect of the graph.** The matrix  $L^\dagger$  is the pseudo-inverse of the Laplacian of the graph, and has a well-known interpretation. Imagine a circuit constructed according to the graph, with  $1\Omega$  resistors along each edge. Let  $R_{ij}$  be the resistance of the resulting circuit between nodes  $i$  and  $j$ . This quantity is called the *resistance distance* of the graph. Then

$$R_{ij} = (e_i - e_j)^\top L^\dagger (e_i - e_j)$$

where  $e_i$  is the canonical basis vector [11, 16]. This interpretation leads to several intuitive consequences. For example, we have *Rayleigh monotonicity*, the fact that adding an edge cannot increase any  $R_{ij}$ . It is also immediately clear that the resistance distance between any two nodes is less than or equal to the path length between those nodes.

Theorems 2 and 3 show that the effect of the graph on the  $L_2$  performance of the bittide system is entirely through the matrix  $L^\dagger$  of resistance distances of the graph. In particular, we can see that *adding edges can only improve performance*. If we add an edge to the graph, then the resistance distance between any pair of nodes cannot increase. Therefore, if the norms of frequency deviation and relative buffer occupancy change, they must decrease.

## 5.2 Two Disequilibrated Frequencies

An illustrative situation for a bittide synchronization system is when the system is almost in equilibrium, except for two nodes. Consider a system in which there are two nodes,  $i$  and  $j$ , that have the frequencies  $1 + \alpha$  and  $1 - \alpha$ , respectively. All other nodes have uncorrected frequency 1. That is

$$\omega^u = \mathbf{1} + \alpha(e_i - e_j)$$

Since  $L^\dagger \mathbf{1} = 0$ , we have

$$\omega^u \top L^\dagger \omega^u = R_{ij}$$

and using this, Theorems 2 and 3 give the performance explicitly in terms of the resistance distance between nodes  $i$  and  $j$  as

$$\|\omega - \omega^{\text{ss}}\|^2 = \frac{\alpha^2 R_{ij}}{2a} \quad \|\delta\|^2 = \frac{\alpha^2 R_{ij}}{2ab}$$

Interpreting these results, we see that the norm performance for both frequency and occupancy scales with the square root of the resistance distance between the nodes.




Figure 2: Frequency for a system with two closely-spaced disequilibrated nodes




Figure 3: Frequency for a system with two distant disequilibrated nodes

For example, consider a bittide system on a  $4 \times 6$  mesh graph, and let  $\alpha = 10^{-4}$ . The parameters of this system are  $k_P = 2 \times 10^{-8}$ ,  $k_I = 10^{-15}$ ,  $l_{ji} = 5000$ ,  $p = 10^5$ ,  $d = 1000$ ,  $\theta_i^0 = 0.1$ ,  $\omega^{(-2)} = \omega^u$ ,  $\omega^{(-1)} = \omega^u$ , and  $\omega^c = 1$ . We consider two examples, one in Figure 2 where nodes  $i$  and  $j$  have small resistance distance  $R \approx 0.700$ , and another in Figure 3 where  $R \approx 2.262$ . In both figures the graph is shown with node  $i$  highlighted in red and node  $j$  highlighted in blue. As a consequence, we have for Figure 2

$$\|\omega - \omega^{\text{ss}}\|^2 \approx 0.175$$

and for Figure 3

$$\|\omega - \omega^{\text{ss}}\|^2 \approx 0.565$$

The greater resistance between the perturbed nodes in Figure 3 leads to worse performance, as shown by the slower convergence of frequency divergence in the figure.

## 5.3 Worse-Case Frequencies

Another application of Theorems 2 and 3 is that they allow computation of the worst-case uncontrolled fre-

quency distribution  $\omega^u$ . The norm response of both occupancy and frequency deviation is proportional to

$$f(\omega^u) = \omega^{u\top} L^\dagger \omega^u$$

We consider all uncontrolled frequencies such that

$$\|\omega^u\|_2 \leq \gamma$$

and seek to maximize  $f(\omega^u)$  over this bounded set. Bounding the set is important for the problem to be mathematically meaningful, since otherwise we can make  $f(\omega^u)$  large simply by scaling  $\omega^u$ . However, the results here are determined also by the way in which we have chosen to bound  $\omega^u$ . Here we choose the Euclidean norm primarily because for this choice we can compute exactly the maximum of  $f$  and the corresponding worst-case  $\omega^u$ . Such a choice could be motivated by assuming a Gaussian probabilistic model for  $\omega^u$ . Alternatively, a deterministic formulation might be better suited to an analysis using the  $\infty$ -norm. We do not delve further into these alternatives here, but instead view the choice of set as a rough proxy for a more accurate model of the set of possible uncontrolled frequencies.

The  $x$  that maximize a homogeneous quadratic function  $x^\top Q x$  over  $x \in \mathbb{R}^n$  with  $\|x\| \leq 1$  is given by  $x = v$ , where  $v$  is the unit eigenvector of  $Q$  with largest eigenvalue. Here we consider  $Q = L^\dagger$ . The second smallest eigenvalue of  $L$  is called the *algebraic connectivity* of the graph, and the corresponding unit eigenvector is called the *Fiedler vector* [7]. Since exactly one eigenvalue of  $L$  is zero, and the others are strictly positive, the eigenvector  $v$  that maximizes  $x^\top Q x$  is the Fiedler vector.




Figure 4: Worst-case uncontrolled frequencies

Figure 4 shows the corresponding worst-case frequency distributions for three example graphs. Here red shows positive values of  $\omega_i^c$  and blue shows negative values; the exact scale is omitted since it is an arbitrary consequence of the magnitude of  $\omega^c$ . Notice that for the rectangular grid graph shown, the worst-case distribution varies from top-to-bottom, whereas for the square grid it varies diagonally, even though the greatest resistance between any two nodes on the rectangular graph is achieved by diagonally opposite corners.

## 6 Conclusions

In this paper we have analyzed the performance of bit-tide synchronization using the 2-norm of the frequency

deviation and relative buffer occupancies. When using PI control, we have shown these quantities are determined in a simple way by the control gains and by the graph resistances. We used these results to analyze and illustrate some simple examples, showing the utility of resistance as a performance indicator in these systems.

## Acknowledgments

We thank Jean-Jacques Slotine for initial guidance on controller behavior. We thank Sahil Hasan and Tong Shen for all their work on the project.

## Appendix: Proofs

Consider the dynamics  $\dot{x} = Ax$ , with output  $y = Cx$  and Lyapunov equation

$$A^\top X + XA + C^\top C = 0$$

We will need the following two standard results from linear systems theory.

**Lemma 4** (Theorem 4.1 in [6]). *If  $A$  is Hurwitz, then the Lyapunov equation has a unique solution  $X \in \mathbb{R}^{n \times n}$  and  $\|y\|_2^2 = x(0)^\top X x(0)$ .*

**Lemma 5** (Proposition 4.2 in [6]). *If  $C^\top C > 0$  and  $X > 0$  satisfies the Lyapunov equation, then  $A$  is Hurwitz.*

First, to reduce the system (9) to this form, we need to remove the non-zero limiting value of the state, which is induced by the constant forcing term  $\omega^u$ . Define  $\bar{x} = \tilde{x} - \tilde{x}^{ss}$  and  $\bar{\omega} = \omega - \omega^{ss}$  then we have dynamics

$$\begin{aligned} \dot{\bar{x}} &= \hat{A}\bar{x} & \bar{x}(0) &= \hat{A}^{-1}\hat{B}_2\omega^u \\ \bar{\omega} &= \hat{C}_1\bar{x} \\ \delta &= \hat{C}_2\bar{x} \end{aligned}$$

Now we can solve the Lyapunov equations. Define

$$X_1 = \begin{bmatrix} \frac{a}{2}\hat{L} + \frac{b}{2a}I & -\frac{b}{2}I \\ -\frac{b}{2}I & \frac{b^2}{2a}\hat{L}^{-1} \end{bmatrix} \quad X_2 = \begin{bmatrix} \frac{1}{2a}I & 0 \\ 0 & \frac{b}{2a}\hat{L}^{-1} \end{bmatrix}$$

**Proof of Theorem 1.** First we show that  $\text{null}(\hat{C}) = \{0\}$ . This follows because

$$\hat{C} = \begin{bmatrix} -aLU_1 & bU_1 \\ -B^\top U_1 & 0 \end{bmatrix}$$

The 1, 2 block satisfies  $\text{null}(U_1) = \{0\}$  since  $U_1$  has orthonormal columns. The 2, 1 block satisfies

$$\text{null}(B^\top U_1) = \text{null}(U_1^\top B B^\top U_1) = \text{null}(\hat{L}) = \{0\}$$

Therefore  $\hat{C}^\top \hat{C} > 0$ . Now one can verify that with  $X = X_1 + X_2$ , we have  $\hat{A}^\top X + X\hat{A} + \hat{C}^\top \hat{C} = 0$ . Finally,  $X_1 > 0$

(via the Schur complement condition) and  $X_2 > 0$ , and hence by Lemma 5 the matrix  $\hat{A}$  is Hurwitz. ■

**Proof of Theorem 2.** One can verify that  $\hat{A}^\top X_1 + X_1 \hat{A} + \hat{C}_1^\top \hat{C}_1 = 0$ . Using Theorem 1 we know that  $\hat{A}$  is Hurwitz, and hence Lemma 4 implies that

$$\begin{aligned} \|\bar{w}\|^2 &= (\hat{A}^{-1} \hat{B}_2 \omega^u)^\top X_1 \hat{A}^{-1} \hat{B}_2 \omega^u \\ &= \frac{1}{2a} \omega^{u\top} L^\dagger \omega^u \end{aligned}$$

as desired. ■

**Proof of Theorem 3.** One can verify that  $\hat{A}^\top X_2 + X_2 \hat{A} + \hat{C}_2^\top \hat{C}_2 = 0$ . Using Theorem 1 we know that  $\hat{A}$  is Hurwitz, and hence Lemma 4 implies that

$$\begin{aligned} \|\delta\|^2 &= (\hat{A}^{-1} \hat{B}_2 \omega^u)^\top X_2 \hat{A}^{-1} \hat{B}_2 \omega^u \\ &= \frac{1}{2ab} \omega^{u\top} L^\dagger \omega^u \end{aligned}$$

as desired. ■

## References

- [1] M. Andreasson, D. V. Dimarogonas, H. Sandberg, and K. H. Johansson. Distributed control of networked dynamical systems: Static feedback, integral action and consensus. *IEEE Transactions on Automatic Control*, 59(7):1750–1764, 2014.
- [2] S. P. Boyd, P. Diaconis, and L. Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667–689, 2004.
- [3] D. Burbano Lombana and M. di Bernardo. Distributed PID control for consensus of homogeneous and heterogeneous networks. *IEEE Transactions on Control of Network Systems*, 2(2):154–163, 2015.
- [4] J. C. Corbett et al. Spanner: Google’s globally distributed database. *ACM Transactions on Computer Systems*, 31(3):1–22, Aug. 2013.
- [5] F. Dorfler and F. Bullo. Synchronization and transient stability in power networks and nonuniform kuramoto oscillators. *SIAM Journal on Control and Optimization*, 50(3):1616–1642, 2012.
- [6] G. E. Dullerud and F. Paganini. *A Course in Robust Control Theory: A Convex Approach*. Springer Verlag, 2000.
- [7] M. Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.
- [8] R. A. Freeman, P. Yang, and K. M. Lynch. Stability and convergence properties of dynamic average consensus estimators. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 338–343, 2006.
- [9] M. Fu and L. Xie. The sector bound approach to quantized feedback control. *IEEE Transactions on Automatic Control*, 50(11):1698–1711, 2005.
- [10] C. Godsil and G. F. Royle. *Algebraic graph theory*. Springer, 2001.
- [11] I. Gutman and W. Xiao. Generalized inverse of the Laplacian matrix and some applications. *Bull. Acad. Serb. Sci. Arts (CI. Math. Natur.)*, pages 15–23, 2004.
- [12] A. Hassibi, S. P. Boyd, and J. P. How. Control of asynchronous dynamical systems with rate constraints on events. In *38th IEEE Conference on Decision and Control*, volume 2, pages 1345–1351, 1999.
- [13] W. K. Hastings. *Monte Carlo sampling methods using Markov chains and their applications*. Oxford University Press, 1970.
- [14] A. Jadbabaie, J. Lin, and S. A. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, 2003.
- [15] F. P. Kelly, A. K. Maulloo, and D. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, 1998.
- [16] D. J. Klein and M. Randić. Resistance distance. *Journal of Mathematical Chemistry*, 12(1):81–95, 1993.
- [17] T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7(1):49–58, 1970.
- [18] S. Lall, C. Caçaval, M. Izzard, and T. Spalink. Modeling and control of bittide synchronization. American Control Conference, 2022. Available at <https://arxiv.org/abs/2109.14111>.
- [19] S. Lall and G. Dullerud. An LMI solution to the robust synthesis problem for multi-rate sampled-data systems. *Automatica*, 37(12):1909–1922, 2001.
- [20] D. L. Mills. Internet time synchronization: the network time protocol. *IEEE Transactions on Communications*, 39(10):1482–1493, 1991.
- [21] L. Mirkin, Z. Palmor, and D. Shneiderman.  $H_2$  optimization for systems with adobe input delays: A loop shifting approach. *Automatica*, 48(8):1722–1728, 2012.
- [22] R. Olfati-Saber and R. Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9):1520–1533, 2004.
- [23] C. W. Reynolds. Flocks, herds and schools: a distributed behavioral model. In *Proceedings 14th ACM SIGGRAPH*, pages 25–34, 1987.
- [24] T. Spalink. *Deterministic sharing of distributed resources*. Princeton University, 2006.
- [25] S. Strogatz. From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1-4):1–20, 2000.
- [26] D. Swaroop and J. K. Hedrick. String stability of interconnected systems. *IEEE Transactions on Automatic Control*, 41(3):349–357, 1996.
- [27] Telcordia GR-253. Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria, 2000.
- [28] W. Wang and J.-J. E. Slotine. On partial contraction analysis for coupled nonlinear oscillators. *Biological Cybernetics*, 92(1):38–53, 2005.
- [29] A. T. Winfree. Biological rhythms and the behavior of populations of coupled oscillators. *Journal of Theoretical Biology*, 16(1):15–42, 1967.