

Skewless Network Clock Synchronization

Enrique Mallada*, Xiaoqiao Meng[†], Michel Hack[†], Li Zhang[†], and Ao Tang*

* School of ECE, Cornell University, Ithaca, NY 14853, USA.

[†] IBM T. J. Watson Research Center, 1101 Kitchawan Rd, Yorktown Heights, NY 10598, USA.

Email: {em464@,atang@ece.}cornell.edu, {xmeng,hack,zhangli}@us.ibm.com

Abstract—This paper examines synchronization of computer clocks connected via a data network and proposes a skewless algorithm to synchronize them. Unlike existing solutions, which either estimate and compensate the frequency difference (skew) among clocks or introduce offset corrections that can generate jitter and possibly even backward jumps, our algorithm achieves synchronization without these problems. We first analyze the convergence property of the algorithm and provide necessary and sufficient conditions on the parameters to guarantee synchronization. We then implement our solution on a cluster of IBM BladeCenter servers running Linux and study its performance. In particular, both analytically and experimentally, we show that our algorithm can converge in the presence of timing loops. This marks a clear contrast with current standards such as NTP and PTP, where timing loops are specifically avoided. Furthermore, timing loops can even be beneficial in our scheme. For example, it is demonstrated that highly connected subnetworks can collectively outperform individual clients when the time source has large jitter. It is also experimentally demonstrated that our algorithm outperforms other well-established software-based solutions such as the NTPv4 and IBM Coordinated Cluster Time (IBM CCT).

I. INTRODUCTION

Keeping consistent time among different nodes in a network is a fundamental requirement of many distributed applications. Their internal clocks are usually not accurate enough and tend to drift apart from each other over time, generating inconsistent time values. Network clock synchronization allows these devices to correct their clocks to match a global reference of time, such as the Universal Coordinated Time (UTC), by performing time measurements through the network. For example, for the Internet, network clock synchronization has been an important subject of research and several different protocols have been proposed [1]–[7]. These protocols are used for various legacy and emerging applications with diverse precision requirements such as banking transactions, communications, traffic measurement and security protection. In particular, in modern wireless cellular networks, time-sharing protocols need an accuracy of several microseconds to guarantee the efficient use of channel capacity. Another example is the recently announced Google Spanner [8], a globally-distributed database, which depends on globally-synchronized clocks within at most several milliseconds drifts.

There are two major difficulties that make the network clock synchronization problem challenging. First, the frequency of hardware clocks is sensitive to temperature and is constantly varying. Second, the latency introduced by the OS and network congestion delay results in errors in the time measurements. Thus, most protocols introduce different ways of estimating the frequency mismatch (skew) [10], [12] and

measuring the time difference (offset) [13], [14]. This leads to extensive literature on skew estimation [10], [15]–[17] which suggests that explicit skew estimation is necessary for clock synchronization.

This paper takes a different approach and shows that focusing on skew estimation could be misleading. We provide a simple algorithm that is able to compensate the clock skew without any explicit estimation of it. Our algorithm only uses current offset information and an exponential average of the past offsets. Thus, it neither needs to store long offset history nor perform expensive computations on them. We analyze the convergence property of the algorithm and provide necessary and sufficient conditions for synchronization. The parameter values that guarantee synchronization depend on the network topology, but there exists a subset of them that is independent of topology and therefore of great practical interest.

We also discover a rather surprising fact. A common practice in the clock synchronization community is to avoid timing loops in the network [1, p. 3] [3, p. 16, s. 6.2]. This is because it is thought that timing loops can introduce instability as stated in [1]: “Drawing from the experience of the telephone industry, which learned such lessons at considerable cost, the subnet topology... must never be allowed to form a loop.” Even though for some parameter values loops can produce instability, we show that a set of proper parameters can guarantee convergence even in the presence of loops. Furthermore, we experimentally demonstrate in Section V that timing loops among clients can actually help reduce the jitter of the synchronization error and is therefore desirable.

A. Related Work and Contribution

Clock synchronization on computer networks has been subject of study for more than 20 years. The current *de facto* standard for IP networks is the Network Time Protocol (NTP) proposed by David Mills [1]. It is a low-cost, purely software-based solution whose accuracy mostly ranges from hundreds of microseconds to several milliseconds, which is often not sufficient. On the other hand, IEEE 1588 (PTP) [3] gives superior performance by achieving sub-microsecond or even nanosecond accuracy. However, it is relatively expensive as it requires special hardware support to achieve those accuracy levels and may not be fully compatible with legacy cluster systems.

More recently, new synchronization protocols have been proposed with the objective of balancing between accuracy and cost. For example, IBM Coordinated Cluster Time (CCT) [11] is able to provide better performance than NTP without additional hardware. Its success is based on a skew estimation mechanism [12] that progressively adapts the clock

frequency without offset corrections. Another alternative is the RADclock [4], [7] which estimates the skew and produces offset corrections, but provides a secondary relative clock that is more robust to jitter.

The solution provided in this paper solves problems present on IBM CCT and RADclock. We are able to achieve microsecond level accuracy without requiring any special hardware as the previous solutions. However, our protocol does not explicitly estimate the skew, which makes the implementation simpler and more robust to jitter than IBM CCT, and does not introduce offset corrections, which avoids the need of a secondary clock as in RADclock. Furthermore, we present a theoretical analysis of its behavior in network environments that unveils some rather surprising facts.

The rest of the paper is organized as follows. In Section II we provide some background on how clocks are actually implemented in computers and how different protocols discipline them. Section III motivates and describes our algorithm together with an intuitive explanation of why it works. In Section IV, we analyze the algorithm and determine the set of parameter values and connectivity patterns under which synchronization is guaranteed. Experimental results evaluating the performance of the algorithm are presented in Section V. We conclude in Section VI.

II. SYNCHRONIZATION OF COMPUTER CLOCKS


Most computer architectures keep their own estimate of time using a counter that is periodically increased by either hardware or kernel's interrupt service routines (ISRs). On Linux platforms for instance, there are usually several different clock devices that can be selected as the clock source by changing the *clocksource* kernel parameter. One particular counter that has recently been used by several clock synchronization protocols [4], [11] is the Time Stamp Counter (TSC) that counts the number of CPU cycles since the last restart of the system. For example, in the IBM BladeCenter LS21 servers, the TSC is a 64-bit counter that increments every $\delta^o = 0.416\text{ns}$ since the CPU nominal frequency $f^o = 1/\delta^o = 2399.711\text{MHz}$.

Based on this counter, each server builds its own estimate $x_i(t)$ of the global time reference, UTC, denoted here by t . Thus, synchronizing computer clocks implies correcting $x_i(t)$ in order to match t , i.e. enforcing $x_i(t) = t$. There are two difficulties on this estimation process. Firstly, the initial time t_0 in which the counter starts its unknown. Secondly, the clock frequency is usually unknown with enough precision and therefore presents a skew $r_i = \frac{x_i(t) - x_i(t_0)}{t - t_0}$. This is illustrated in Figure 1a where $x_i(t)$ not only increases at a different rate than t , but also starts from a value different from t_0 , represented by x_i^o .

Mathematically, $x_i(t)$ can be described by the linear map of the global reference t , i.e.

$$x_i(t) = r_i s_i^o (t - t_0) + x_i^o, \quad (1)$$

where s_i^o is an additional skew correction implemented to compensate the skew r_i ; in Figure 1a $s_i^o = 1$. Equation (1) also shows that if one can set $s_i^o = 1/r_i$ and $x_i^o = t_0$, then we obtain a perfectly synchronized clock with $x_i(t) = t$.




(a) Illustration of computer time estimate $x_i(t)$ and UTC time t (b) Offset and relative skew measurements

Fig. 1: Time estimation and relative measurements

The main problem is that not only neither t_0 nor r_i can be explicitly estimated, but also r_i varies with time as shown in Figure 2a. Thus, current protocols periodically update s_i^o and x_i^o in order to keep track of the changes of r_i . These updates are made using the *offset* between the current estimate $x_i(t)$ and the global time t , i.e. $D_i^x(t) = t - x_i(t)$, and the *relative frequency error* that is computed using two offset measurements separated by τ seconds, i.e.

$$f_i^{err}(t) := \frac{D_i^x(t) - D_i^x(t - \tau)}{x_i(t) - x_i(t - \tau)} = \frac{1 - r_i s_i^o}{r_i s_i^o}. \quad (2)$$

Figure 1b provides an illustration of these measurements.



(a) Offset between two TSC counters: (b) Example of skew and offset corrections using *adjtimex()*. The straight line is a linear fit that is recomputed on linux time: First a $20\mu\text{s}$ is subtracted from the offset values in offset is added and subtracted and residual offset axis then a skew of 0.3ppm is introduced

Fig. 2: Comparison between two TSC counters, and skew and offset corrections using *adjtimex()*

To understand the differences between current protocols, we first rewrite the evolution of $x_i(t)$ based only on the time instants t_k in which the clock corrections are performed. We allow the skew correction s_i^o to vary over time, i.e. $s_i(t_k)$, and write $x_i(t_{k+1})$ as a function of $x_i(t_k)$. Thus, we obtain

$$x_i(t_{k+1}) = x_i(t_k) + \tau r_i s_i(t_k) + u_i^x(t_k) \quad (3a)$$

$$s_i(t_{k+1}) = s_i(t_k) + u_i^s(t_k) \quad (3b)$$

where $\tau = t_{k+1} - t_k$ is the time elapsed between adaptations; also known as poll interval [1]. The values $u_i^x(t_k)$ and $u_i^s(t_k)$ represent two different types of corrections that a given protocol chooses to do at time t_k and are usually implemented within the interval (t_k, t_{k+1}) . $u_i^x(t_k)$ is usually referred to as *offset correction* and $u_i^s(t_k)$ as *skew correction*.¹ See Figure 2b for an illustration of their effect on the linux time.

¹These corrections can be implemented in Linux OS using the *adjtimex()* interface to update the system clock or by maintaining a virtual version of $x_i(t)$ and directly applying the corrections to it, as in IBM CCT [11] and RADclock [4]. The latter gives more control on how the corrections are implemented since it does not depend on kernel's routines.




Fig. 3: Variations of NTP time using TSC as reference




Fig. 4: Current Protocols Adaptation

We now proceed to summarize the different types of adaptations implemented by current protocols. The main differences between them are whether they use offset corrections, skew corrections, or both, and whether they update using offset values $D_i^x(t_k)$, relative frequency errors $f_i^{err}(t_k)$, or both.

A. Offset corrections

These corrections consist in keeping the skew fixed and periodically introducing time changes of size $u_i^x(t_k) = \kappa_1 D_i^x(t_k)$ or $u_i^x(t_k) = \kappa_1 D_i^x(t_k) + \kappa_2 f_i^{err}(t_k)$ where $\kappa_1, \kappa_2 > 0$. They are used by NTPv3 [18] and NTPv4 [1] respectively under ordinary conditions.

These protocols have in general a slow initialization period as shown in Figure 3a. This is because the algorithm must first obtain a very accurate estimate of the initial frequency error $f_i^{err}(t_0)$. Furthermore, these updates usually generate non-smooth time evolutions as in Figures 3b and 4a, and should be done carefully since they might introduce backward jumps ($x_i(t_{k+1}) < x_i(t_k)$), which can be problematic for some applications.

B. Skew corrections

Another alternative that avoids using steep changes in time is proposed by the IBM CCT solution [11]. This alternative does not introduce any offset correction, i.e. $u_i^x(t_k) = 0$, and updates the skew $s_i(t_k)$ by $u_i^s(t_k) = \kappa_1 D_i^x(t_k) + \kappa_2 f_i^{err}(t_k)$.

The behavior of this algorithm is shown in Figure 4b. In [19] it was shown for a slightly modified version of it (used $r_i s_i(t_k) f_i^{err}(t_k)$ instead of $f_i^{err}(t_k)$) the algorithm can achieve synchronization for very diverse network architectures.

However, the estimation of $f_i^{err}(t_k)$ is nontrivial as it is constantly changing with subsequent updates of $s_i(t_k)$ and it usually involves sophisticated computations [10], [12].




Fig. 5: Unstable clock steering using only offset information (4) and stable clock steering based on exponential average compensation (6)

C. Skew and offset corrections

This type of corrections allow dependence on only offset information $D_i^x(t_k)$ as input to $u_i^x(t_k)$ and $u_i^s(t_k)$. For instance, in [5] the update $u_i^x(t_k) = \kappa_1 D_i^x(t_k)$ and $u_i^s(t_k) = \kappa_2 D_i^x(t_k)$ was proposed.

This option allows the system to achieve synchronization without any skew estimation. But the cost of achieving it, is introducing offset corrections in $x_i(t)$ as shown in Figure 4c. Therefore, it suffers from the same problems discussed in II-A.

III. SKEWLESS NETWORK SYNCHRONIZATION

We now present an algorithm that overcomes the limitations of the solutions described in Section II. In other words, our solution has the following two properties:

- 1) Continuity: The protocol does not introduce steep changes on the time value, i.e. $u_i^x(t_k) \equiv 0$.
- 2) Skew independence: The protocol does not use skew information $f_i^{err}(t_k)$ as input.

A solution with these properties will therefore prevent unnecessary offset corrections that produce jitter and will be more robust to noise by avoiding skew estimation. After describing and motivating our algorithm, we show how the updating rule can be implemented in the context of a network environment.

The motivation behind the proposed solution comes from trying to compensate the problem that arises when one tries to naively impose properties 1) and 2), i.e. using

$$u_i^x(t_k) = 0 \quad \text{and} \quad u_i^s(t_k) = \kappa_1 D_i^x(t_k). \quad (4)$$

Figure 5 shows that this type of clock corrections is unstable; the offset $D_i^x(t_k)$ of the slave clock oscillates with an exponentially increasing amplitude.

The oscillations in Figure 5 arise due to the fundamental limitations of using offset to update frequency. This is better seen in the continuous time version of the system (3) with (4), i.e.

$$\dot{x}_i(t) = r_i s_i(t) \quad \text{and} \quad \dot{s}_i(t) = \kappa_1 D_i^x(t)$$

where $\dot{x}(t) = \frac{d}{dt}x(t)$. If we consider the offset $D_i^x = t - x_i(t)$ as the system state, then we have

$$\dot{D}_i^x = 1 - r_i s_i \quad \text{and} \quad \ddot{D}_i^x = -\kappa_1 r_i D_i^x,$$

with $\ddot{x}(t) = \frac{d^2}{dt^2}x(t)$.

This is analogous to a spring mass system without friction. Thus, it has two purely imaginary eigenvalues that generate

sustained oscillations; see [6], [20] for similar examples.² One way to damp these oscillations in the spring-mass case is by adding *friction*. This implies adding a term that includes a frequency mismatch $f_i^{err}(t)$ in our system, which is equivalent to the protocols of Section II-B, and therefore undesired.

However, there are other ways to damp these oscillations using passivity-based techniques from control theory [21], [22]. The basic idea is to introduce an additional state y_i that generates the desired *friction* to damp the oscillations.

Inspired by [21], we consider the exponentially weighted moving average of the offset

$$y_i(t_{k+1}) = pD_i^x(t_k) + (1-p)y_i(t_k). \quad (5)$$

and update $x_i(t_k)$ and $s_i(t_k)$ using:

$$u_i^x(t_k) = 0 \quad \text{and} \quad u_i^s(t_k) = \kappa_1 D^x(t_k) - \kappa_2 y(t_k). \quad (6)$$

Figure 5 shows how the proposed strategy is able to compensate the oscillations without needing to estimate the value of $f_i^{err}(t_k)$. The stability of the algorithm will depend on how κ_1 , κ_2 and p are chosen. A detailed specification of these values is given in Section IV-B.

Finally, since we are interested in studying the effect of timing loops, we move away from the client-server configuration implicitly assumed in Section II and allow mutual or cyclic interactions among nodes. The interactions between different nodes is described by a graph $G(V, E)$, where V represents the set of n nodes ($i \in V$) and E the set of *directed* edges ij ; $ij \in E$ means node i can measure its offset with respect to j , $D_{ij}^x(t_k) = x_j(t_k) - x_i(t_k)$.

Within this context, a natural extension of (5)-(6) is to substitute $D_i^x(t_k)$ with the weighted average of i 's neighbors offsets. Thus, we propose the following algorithm to update the clocks in the network.

Algorithm 1 (Alg1): For each computer node i in the network, perform the following actions:

- Compute the time offsets ($D_{ij}^x(t_k)$) from i to every neighbor j at time t_k .
- Update the skew $s_i(t_{k+1})$ and the moving average $y_i(t_{k+1})$ at time t_{k+1} according to:

$$s_i(t_{k+1}) = s_i(t_k) + \kappa_1 \sum_{j \in \mathcal{N}_i} \alpha_{ij} D_{ij}^x(t_k) - \kappa_2 y_i(t_k) \quad (7a)$$

$$y_i(t_{k+1}) = p \sum_{j \in \mathcal{N}_i} \alpha_{ij} D_{ij}^x(t_k) + (1-p)y_i(t_k) \quad (7b)$$

where \mathcal{N}_i represents the set of neighbors of i and the weights α_{ij} are positive.

Using this algorithm, many servers can affect the final frequency of the system. Thus, when the system synchronizes, we have

$$x_i(t_k) = r^*(t_k - t_0) + x^* \quad i \in V. \quad (8)$$

r^* and x^* are possibly different from their ideal values 1 and t_0 . Their final values depend on the initial condition of all different clocks as well as the topology, which we assume to be a connected graph in this paper.

²In the discrete time system the oscillations increase in amplitude since there is a delay between the time the offset is measured t_k and the time the update is made t_{k+1} which makes the system unstable.

IV. ANALYSIS

We now analyze the asymptotic behavior of system (7) and provide a necessary and sufficient condition on the parameter values that guarantee its convergence to (8). The techniques used are drawn from the control literature, e.g. [5] and [19], yet its application in our case is nontrivial.

Notation: We use $\mathbf{0}_{m \times n}$ ($\mathbf{1}_{m \times n}$) to denote the matrices of all zeros (ones) within $\mathbb{R}^{m \times n}$ and $\mathbf{0}_n$ ($\mathbf{1}_n$) to denote the column vectors of appropriate dimensions. $I_n \in \mathbb{R}^{n \times n}$ represents the identity matrix. Given a matrix $A \in \mathbb{R}^{n \times n}$ with Jordan normal form $A = PJP^{-1}$, let $n_A \leq n$ denote the total number of Jordan blocks J_l with $l \in \mathcal{I}(A) := \{1, \dots, n_A\}$. We use $\mu_l(A)$, $l \in \{1, \dots, n\}$ or just $\mu(A)$ to denote the eigenvalues of A , and order them decreasingly $|\mu_1(A)| \geq \dots \geq |\mu_n(A)|$. Finally, A^T is the transpose of A , A_{ij} is the element of the i th row and j th column of A and a_i is the i th element of the column vector a , i.e. $a = [a_i]^T$.

It is more convenient for the analysis to use a vector form representation of (7) given by

$$z_{k+1} = Az_k \quad (9)$$

where $z_k := [x(t_k)^T s(t_k)^T y(t_k)^T]^T \in \mathbb{R}^{3n}$,

$$A := \begin{bmatrix} I_n & \tau R & \mathbf{0}_{n \times n} \\ -\kappa_1 L & I_n & -\kappa_2 I_n \\ p(-L) & \mathbf{0}_{n \times n} & (1-p)I_n \end{bmatrix} \in \mathbb{R}^{3n \times 3n},$$

$R \in \mathbb{R}^{n \times n}$ is the diagonal matrix with elements r_i and $L \in \mathbb{R}^{n \times n}$ is the Laplacian matrix associated with $G(V, E)$,

$$L_{ii} = \alpha_{ii} := \sum_{j \in \mathcal{N}_i} \alpha_{ij} \quad \text{and} \quad L_{ij} = \begin{cases} -\alpha_{ij} & \text{if } ij \in E, \\ 0 & \text{otherwise.} \end{cases}$$

The convergence analysis of this section is done in two stages. First, we provide necessary and sufficient conditions for synchronization in terms of the eigenvalues of A (Section IV-A) and then use Hermite-Biehler Theorem [23] to relate these eigenvalues with the parameter values that can be directly used in practice (Section IV-B). All the proof details are included in the appendix for interested readers.

A. Asymptotic Behavior

We start by studying the asymptotic behavior of (9). That is, we are interested in finding under what conditions the series of elements $\{x_i(t_k)\}$ converge to (8) as t_k goes to infinity.

Consider the Jordan normal form [24] of

$$A := [\zeta_1 \quad \dots \quad \zeta_{3n}] J [\eta_1 \quad \dots \quad \eta_{3n}]^T$$

where $J = \text{blockdiag}(J_l)_{l \in \mathcal{I}(A)}$, ζ_i and η_i are the right and left generalized eigenvectors of A such that

$$\zeta_i^T \eta_j = \begin{cases} 1 & \text{if } j = i, \\ 0 & \text{otherwise.} \end{cases}$$

The crux of the analysis comes from understanding the relationship between the multiplicity of the eigenvalue $\mu(A) = 1$ and the eigenvalue $\mu(L) = 0$, and their corresponding eigenvectors. This is captured in the next two lemmas.

Lemma 1 (Eigenvalues of A and Multiplicity of $\mu(A) = 1$): A has an eigenvalue $\mu(A) = 1$ with multiplicity 2 if and only if the graph $G(V, E)$ is connected, $\kappa_1 \neq \kappa_2$ and $p > 0$.

Furthermore, $\mu_l(A)$ are the roots of

$$g_l(\lambda) := (\lambda - 1)^2(\lambda - 1 + p) + [(\lambda - 1)\kappa_1 + \kappa_2 - \kappa_1]\nu_l \quad (10)$$

where $\nu_l = \mu_l(\tau LR)$ and satisfies

$$\nu_n = 0 < |\nu_l| \text{ for } l \in \{1, \dots, n-1\}. \quad (11)$$

Lemma 2 (Jordan Chains of $\mu(A) = 1$ and $\mu(A) = 1 - p$):

Under the conditions of Lemma 1 the right and left Jordan chains, (ζ_1, ζ_2) and (η_2, η_1) respectively, associated with $\mu(A) = 1$ and the eigenvectors ζ_3 and η_3 associated with $\mu(A) = 1 - p$ are given by

$$[\zeta_1 \ \zeta_2 \ \zeta_3] = \begin{bmatrix} \mathbf{1}_n & \mathbf{1}_n & -\frac{\tau\kappa_2}{p^2}\mathbf{1}_n \\ \mathbf{0}_n & \frac{(R^{-1}\mathbf{1}_n)}{\tau} & \frac{\kappa_2}{p}R^{-1}\mathbf{1}_n \\ \mathbf{0}_n & \mathbf{0}_n & R^{-1}\mathbf{1}_n \end{bmatrix} \text{ and} \quad (12)$$

$$[\eta_1 \ \eta_2 \ \eta_3] = \gamma \begin{bmatrix} R^{-1}\xi & \mathbf{0}_n & \mathbf{0}_n \\ -\tau\xi & \xi & \mathbf{0}_n \\ \tau\kappa_2(\frac{1}{p} + \frac{1}{p^2})\xi & -\frac{\kappa_2}{p}\xi & \xi \end{bmatrix} \quad (13)$$

where ξ is the unique normalized left eigenvector of $\mu(L) = 0$ ($\sum_{i=1}^n \xi_i = 1$) and γ is the ξ_i -weighted harmonic mean of r_i , i.e. $\frac{1}{\gamma} = \mathbf{1}_n^T R^{-1}\xi = \sum_{i=1}^n \frac{\xi_i}{r_i}$.

The proof of Lemmas 1 and 2 can be found in the Appendices A and B. We now proceed to state our main convergence result.

Theorem 1 (Convergence): The algorithm (9) achieves synchronization for any initial conditions if and only if the graph $G(V, E)$ is connected, $\kappa_1 \neq \kappa_2$, $p > 0$ and $|\mu_l(A)| < 1$ whenever $\mu_l(A) \neq 1$. Moreover, whenever the system synchronizes, we have

$$x^* = \gamma \sum_{i=1}^n \xi_i \left(\frac{1}{r_i} x_i(t_0) + \tau \frac{\kappa_2}{p^2} y_i(t_0) \right), \text{ and} \quad (14a)$$

$$r^* = \gamma \sum_{i=1}^n \xi_i (s_i(t_0) - \frac{\kappa_2}{p} y_i(t_0)). \quad (14b)$$

Theorem 1 provides an analytical tool to understand the influence of the different nodes of the graph in the final offset x^* and frequency r^* . For example, suppose that we know that node 1 has perfect knowledge of its own frequency (r_1) and the UTC time at $t = t_0$ ($x_1(t_0) = t_0$), and configure the network such that node 1 is the **unique leader** like the top node in Figures 6a and 6c. It is easy to show that $\xi_1 = 1$ and $\xi_i = 0 \ \forall i \neq 1$. Then, using (14a)-(14b) and definition of γ we can see that $\gamma = r_1$ and

$$x^* = x_1(t_0) + r_1 \tau \frac{\kappa_2}{p^2} y_1(t_0) \text{ and } r^* = r_1 s_1(t_0) - \frac{r_1 \kappa_2}{p} y_1(t_0).$$

However, since node 1 knows r_1 and t_0 , it can choose $x_1(t_0) = t_0$, $s_1(t_0) = \frac{1}{r_1}$ and $y_1(t_0) = 0$. Thus, we obtain $x^* = t_0$ and $r^* = 1$ which implies by (8) that every node in the network will end up with $x_i(t) = t$. In other words, Theorem 1 allows us to understand how the information propagates and how we can guarantee that every server will converge to the desired time. Notice that the initial condition used for server 1 is equivalent to assuming that server 1 is a reliable source of UTC like an atomic clock for instance.




Fig. 6: Graphs with real eigenvalue Laplacians

B. Necessary and sufficient conditions for synchronization

We now provide necessary and sufficient conditions in terms of explicit parameter values (κ_1 , κ_2 , τ and p) for Theorem 1 to hold. We will restrict our attention to graphs that have Laplacian matrices with real eigenvalues. This includes for example trees (Figure 6a), symmetric graphs with $\alpha_{ij} = \alpha_{ji}$ (Figure 6b) and symmetric graphs with a leader (Figure 6c).

The proof consists on studying the Schur stability of $g_l(\lambda)$ and has several steps. We first perform a change of variable that maps the unit circle onto the left half-plane. This transforms the problem of studying the Schur stability into a Hurwitz stability problem which is solved using Hermite-Biehler Theorem.

Theorem 2 (Hurwitz Stability (Hermite-Biehler) [23]):

Given the polynomial $P(s) = a_n s^n + \dots + a_0$, let $P^r(\omega)$ and $P^i(\omega)$ be the real and imaginary part of $P(j\omega)$, i.e. $P(j\omega) = P^r(\omega) + jP^i(\omega)$. Then $P(s)$ is a Hurwitz polynomial if and only if

- 1) $a_n a_{n-1} > 0$ and 2)
- 2) The zeros of $P^r(\omega)$ and $P^i(\omega)$ are all simple and real and interlace as ω runs from $-\infty$ to $+\infty$.

We now determine the proper parameter values that guarantee synchronization.

Theorem 3 (Parameter Values for Synchronization):

Given a connected graph $G(V, E)$ such that the corresponding Laplacian matrix L has real eigenvalues. The system (9) achieves synchronization if and only if

- (i) $|1 - p| < 1$ or equivalently $2 > p > 0$
- (ii) $\frac{2\kappa_1}{3p} > \kappa_1 - \kappa_2 > 0$ and (iii) $\tau < \frac{p(\kappa_2 - p(\kappa_1 - \kappa_2))}{\mu_{\max}(\kappa_1 - p(\kappa_1 - \kappa_2))^2}$

where μ_{\max} is the largest eigenvalue of LR .

Even though μ_{\max} depends on r_i which is in general unknown, it is easy to show that $\mu_l(LR) \leq \hat{r}_{\max} \mu_l(L)$ where \hat{r}_{\max} is an upper bound of the maximum rate deviation r_i . Furthermore, using Greshgorin's circle theorem, it is easy to show that $\mu_{\max}(L) \leq 2\alpha_{\max} := 2 \max_i \alpha_{ii}$. Therefore, if we set

$$\tau < \frac{p(\kappa_2 - \delta\kappa p)}{2\alpha_{\max} \hat{r}_{\max} (\kappa_1 - \delta\kappa p)^2} \quad (15)$$

convergence is guaranteed for **every connected graph** with real eigenvalues.

V. EXPERIMENTS

To test our solution and analysis, we implement an asynchronous version of Algorithm 1 (Alg1) in C using the IBM




Fig. 7: Testbed of IBM BladeCenter blade servers

CCT solution as our code base. Every node perform its own measurements and updates every τ seconds using (7), but not necessarily at the same instants t_k .

Our program reads the TSC counter directly using the `rdtsc` assembly instruction to minimize reading latencies and maintains a virtual clock that can be directly updated. The list of neighbors is read from a configuration file and whenever there is no neighbor, the program follows the local Linux clock. Finally, offset measurements are taken using an improved ping pong mechanism proposed in [11].

We run our skewless protocol in a cluster of IBM BladeCenter LS21 servers with two AMD Opteron processors of 2.40GHz, and 16GB of memory. As shown in Figure 7, the servers `serv1-serv10` are used to run the protocol. The offset measurements are taken through a Gigabit Ethernet switch. Server `serv0` is used as a reference node and gathers time information from the different nodes using a Cisco 4x InfiniBand Switch that supports up to 10Gbps between any two ports and up to 240Gbps of aggregate bandwidth. This minimizes the error induced by the data collecting process.

We use this testbed to validate the analysis in Section IV. First, we illustrate the effect of different parameters and analyze the effect of the network configuration on convergence (Experiment 1). Then we present a series of configurations that demonstrate how connectivity between clients is useful in reducing the jitter of a noisy clock source (Experiment 2). And finally, we compare the performance of our protocol with respect to NTP version 4 (Experiment 3) and IBM CCT (Experiment 4).

We will use several performance metrics to evaluate Alg1. For instance, the *mean relative deviation* from the leader which is defined as the root mean square of the node's offset with respect to the leader, i.e. $\sqrt{S_n}$ with

$$S_n = \frac{1}{n-1} \sum_{i=2}^n \langle (x_i - x_1)^2 \rangle, \quad (16)$$

where $\langle \cdot \rangle$ amounts to the sample average. We will also use the 99% Confidence Interval CI_{99} and the maximum offset (CI_{100}) as metrics of accuracy. For example, if $CI_{99} = 10\mu s$, then 99% of the offset samples will be within $10\mu s$ of the leader.

Unless explicitly stated, the default parameter values are

$$p = 0.99, \quad \kappa_1 = 1.1, \quad \kappa_2 = 1.0 \text{ and } \alpha_{ij} = \frac{c}{|\mathcal{N}_i|}. \quad (17)$$

The scalar c is a commit or gain factor that will allow us to compensate the effect of τ . Notice that by definition of α_{ij} , $\alpha_{ii} = c$ for every node that is not the leader.

Moreover, these values immediately satisfy (i) and (ii) of Theorem 3 since $1-p = 0.01$ and $\frac{2\kappa_1}{3p} = 0.7407 > \kappa_1 - \kappa_2 = 0.1$. The remaining condition can be satisfied by modifying τ or equivalently c . Here, we choose to fix $c = 0.7$ which makes condition (iii)

$$\tau < \frac{890.1}{\mu_{\max}} \text{ms.}$$

For fixed polling interval τ , the stability of the system depends on the value of μ_{\max} , which is determined by the underlying network topology and the values of α_{ij} .





Fig. 8: Effect of topology on convergence: (a) Client-server configuration; (b) Two clients connected to server and mutually connected.

Experiment 1 (Convergence): We first consider the client server configuration described in Figure 8a with a time step $\tau = 1s$. In this configuration $\mu_{\max} \approx c = 0.7$ and therefore condition (iii) becomes $\tau < 1.2717s$. Figure 9a shows the offset between `serv1` (the leader) and `serv2` (the client) in microseconds. There we can see how `serv2` gradually updates $s_2(t_k)$ until the offset becomes insignificant.



(a) Client server configuration with $\tau = 1s$. The client converges and the algorithm is stable. (b) Two clients mutually connected with $\tau = 1s$. The algorithm becomes unstable.

Fig. 9: Loss of stability by change in the network topology

Figure 9a tends to suggest that the set of parameters given by (17) and $\tau = 1s$ are suitable for deployment on the servers. This is in fact true provided that network is a directed tree as in Figure 6a. The intuition behind this fact is that in a tree, each client connects only to one server. Thus, those connected to the leader will synchronize first and then subsequent layers will follow.

However, once loops appear in the network, there is no longer a clear dependency since two given nodes can mutually get information from each other. This type of dependency might make the algorithm unstable. Figure 9b shows an experiment with the same configuration as Figure 9a in which `serv2` synchronizes with `serv1` until a third server (`serv3`) appears after 60s. At that moment the system is reconfigured to have the topology of Figure 8b introducing a timing loop

between serv2 and serv3. This timing loop makes the system unstable.

The instability arises since after serv3 starts, the new topology has $\mu_{\max} \approx 1.5c = 1.05$. Thus, the time step condition (iii) becomes $\tau < 847.8\text{ms}$ which is no longer satisfied by $\tau = 1\text{s}$.

This may be solved for the new topology (Figure 8b) by using any τ smaller than 847.8ms. However, if we want a set of parameters that is independent of the topology, we can use (15) and notice that $\alpha_{\max} = c$ and $\hat{r}_{\max} \approx 1$. We choose

$$\tau = 500\text{ms} < \frac{890.2}{2\alpha_{\max}}\text{ms} = \frac{890.2}{2c}\text{ms} = 635.9\text{ms}.$$

Figure 10 shows how now serv2 and serv3 can synchronize with serv1 after introducing this change.




Fig. 10: Two clients mutually connected with $\tau = 500\text{ms}$

Experiment 2 (Timing Loops Effect): We now show how timing loops can be used to collectively outperform individual clients when the time source is noisy.

We run Alg1 on 10 servers (serv1 through serv10). The connection setup is described in Figure 11. Every node is directly connected unidirectionally to the leader (serv1) and bidirectionally to $2K$ additional neighbors.




Fig. 11: Leader topologies with $2K$ neighbors connection. Connections to the leader (serv1) are unidirectional while the connections among clients (serv2 through serv10) are bidirectional

When $K = 0$ then the network reduces to a star topology and when $K = 4$ the servers serv2 through serv10 form a complete graph.

The dashed arrows in Figure 11 show the connections where jitter was introduced. To emulate a link with jitter we added random noise η with values taken uniformly from $\{0, 1, \dots, \text{Jitter}_{\max}\}$ on both direction of the communication,

$$\eta \in \{0, 1, \dots, \text{Jitter}_{\max}\}\text{ms}. \quad (18)$$

Notice that the arrow only shows a dependency relationship, the ping pong mechanism sends packets in both

direction of the physical communication. We used a value of $\text{Jitter}_{\max} = 10\text{ms}$. Since the error was introduced in both directions of the ping pong, this is equivalent to a standard deviation of 6.05ms.




Fig. 12: Offset of the nine servers connected to a noisy clock source

Figure 12 illustrates the relative offset between the two extreme cases; The star topology ($K = 0$) is shown in Figure 12a, and the complete subgraph ($K = 4$) is shown in Figure 12b.

The worst case offset for $K = 0$ is $CI_{100} = 5.1\text{ms}$ which is on the order of the standard deviation of the jitter. However, when $K = 4$ we obtain a worst case offset of $CI_{100} = 690.8\mu\text{s}$, an order of magnitude improvement.




Fig. 13: Effect of the client's communication topology on the mean relative deviation. As the connectivity increases (K increases) the mean relative deviation is reduced by factor of 6.26, i.e. a noise reduction of approx. 8dB.


The change on the mean relative deviation $\sqrt{S_n}$ as the connectivity among clients increases from isolated nodes ($K = 0$) to a complete subgraph ($K = 4$) is studied in Figure 13. The results presented show that even without any offset filtering mechanism the network itself is able to perform a distributed filtering that achieves an improvement of up to a factor of 6.26 or equivalently a noise reduction of almost 8dB.

Experiment 3 (Comparison with NTPv4): We now perform a thorough comparison between our protocol (Alg1) and NTPv4. We will use the one hop configuration of Figure 8b but without the bidirectional link. Here, server serv1 is set as NTP server and as leader of Alg1, server serv2 has a client running NTPv4 and server serv3 a client running our protocol.

In order to make a fair comparison, we need both algorithms to use the same polling interval. Thus, we fix $\tau = 16\text{sec}$. This can be done for NTP by setting the parameters `minpoll` and `maxpoll` to 4 ($2^4 = 16\text{secs}$). The remainder parameter values for Alg1 are given by

$$p = 1.98, \quad \kappa_1 = 1.388 \text{ and } \kappa_2 = 1.374. \quad (19)$$

Figure 14a shows the time differences between the clients running NTPv4 and Alg1 (serv2 and serv3), and the leader (serv1) over a period of 30 hours. It can be seen that Alg1 is able to track serv1’s clock keeping an offset smaller than $10\mu\text{s}$ for most of the time while NTPv4 incurs in larger offsets during the same period of time. This difference is produced by the fact that Alg1 is able to react more rapidly to frequency changes while NTPv4 incurs in more offset corrections that generate larger jitter.



(a) Offset values of NTPv4 and Alg1 (b) Cumulative Distribution Function for a period of 30 hours.

Fig. 14: Performance evaluation between our solution (Alg1) and NTPv4

A more detailed and comprehensive analysis is presented in Figure 14b where we plot the Cumulative Distribution Function (CDF) of the offset samples. That is, the fraction of samples whose time offset is smaller than a specific value. Using Figure 14b we compute the corresponding 99% confidence intervals (CI_{99})

Alg1 achieves a performance of $\sqrt{S_n} = 3.1\mu\text{s}$, $CI_{99} = 9.5\mu\text{s}$ and a maximum offset of $CI_{100} = 15.9\mu\text{s}$, while NTPv4 obtains $\sqrt{S_n} = 8.1\mu\text{s}$, $CI_{99} = 21.8\mu\text{s}$ and a maximum offset of $CI_{100} = 28.0\mu\text{s}$. Thus, not only Alg1 achieves a reduction of $\sqrt{S_n}$ by a factor of 2.6 (-4.2dB) with respect to NTPv4, but it also obtains smaller confidence intervals and maximum offset values.





Fig. 15: Offset values of NTPv4 and Alg1 after a 25ms offset introduced in serv1.

Finally, we investigate the speed of convergence. Starting from both clients synchronized to server serv1, we introduce a 25ms offset. Figure 15 shows how Alg1 is able to converge to a $20\mu\text{s}$ range within one hour while NTPv4 needs 4.5hours to achieve the same synchronization precision.

Experiment 4 (Comparison with IBM CCT): We now proceed to compare the performance of Alg1 with respect to IBM CCT. Notice that unlike IBM CCT, our solution does not perform any previous filtering of the offset samples, the filtering is performed instead by calibrating the parameters which mostly depend on the polling interval τ chosen. Here we use $c = 0.70$, $\tau = 250\text{ms}$, $\kappa_1 = 0.1385$, $\kappa_2 = 0.1363$ and $p = 0.62$.

In Figure 16a we present the mean relative deviation $\sqrt{S_n}$ for two clients connected directly to the leader as the jitter is increased from $\text{Jitter}_{\text{max}} = 0\mu\text{s}$ (no jitter) to $\text{Jitter}_{\text{max}} = 160\mu\text{s}$ with a granularity of $1\mu\text{s}$. The worst case offset is shown in Figure 16b. Each data point is computed using a sample run of 250 seconds.

Our algorithm consistently outperforms IBM CCT in terms of both $\sqrt{S_n}$ and worst case offset. The performance improvement is due to two reasons. Firstly, the noise filter used by the IBM CCT algorithm is tailored for noise distributions that are mostly concentrated close to zero with sporadic large errors. However, it does not work properly in cases where the distribution is more homogeneous as in this case. Secondly, by choosing $\delta\kappa = \kappa_1 - \kappa_2 = 0.002 \ll 1$ the protocol becomes very robust to offset errors.



(a) Mean relative deviation $\sqrt{S_n}$ (b) Maximum offset

Fig. 16: Performance evaluation between our solution (Alg1) and IBM CCT

VI. CONCLUSION

This paper presents a clock synchronization protocol that is able to synchronize networked nodes without explicit estimation of the clock skews and steep corrections on the time. Unlike current standards, our protocol is guaranteed to converge even in the presence of timing loops which allows different clients to share timing information and even collectively outperform individual clients. We implemented our solution on a cluster of IBM BladeCenter servers and empirically verified our predictions and our protocol’s supremacy over several existing solutions.

REFERENCES

- [1] D. Mills, “Network time protocol version 4 reference and implementation guide,” University of Delaware, Tech. Rep. 06-06-1, Jun. 2006.
- [2] A. Sobeih *et al.*, “Almost peer-to-peer clock synchronization,” *Parallel and Distributed Processing Symposium, International*, p. 21, 2007.
- [3] “IEEE standard for a precision clock synchronization protocol for networked measurement and control systems,” pp. 1–269, 2008.
- [4] D. Veitch, J. Ridoux, and S. B. Korada, “Robust synchronization of absolute and difference clocks over networks,” *IEEE/ACM Trans. Netw.*, vol. 17, no. 2, pp. 417–430, Apr. 2009.
- [5] R. Carli and S. Zampieri, “Networked clock synchronization based on second order linear consensus algorithms,” in *Decision and Control (CDC), 2010 49th IEEE Conference on*, Dec. 2010, pp. 7259–7264.
- [6] E. Mallada and A. Tang, “Distributed clock synchronization: Joint frequency and phase consensus,” in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, Dec. 2011, pp. 6742–6747.
- [7] J. Ridoux, D. Veitch, and T. Broomhead, “The case for feed-forward clock synchronization,” *IEEE/ACM Transactions on Networking*, vol. 20, no. 1, pp. 231–242, 2012.
- [8] J. C. Corbett *et al.*, “Spanner: Google’s globally-distributed database,” in *Proceedings of the 10th USENIX conference on Operating Systems Design and Implementation*, ser. OSDI’12. Berkeley, CA, USA: USENIX Association, 2012, pp. 251–264. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2387880.2387905>

- [9] B. Sundararaman, U. Buy, and A. D. Kshemkalyani, "Clock synchronization for wireless sensor networks: a survey," *Ad Hoc Networks*, vol. 3, no. 3, pp. 281–323, 2005.
- [10] H. Kim, X. Ma, and B. Hamilton, "Tracking low-precision clocks with time-varying drifts using kalman filtering," *Networking, IEEE/ACM Transactions on*, vol. 20, no. 1, pp. 257–270, Feb. 2012.
- [11] S. Froehlich *et al.*, "Achieving precise coordinated cluster time in a cluster environment," in *Precision Clock Synchronization for Measurement, Control and Communication, 2008. ISPCS 2008. IEEE International Symposium on*, Sep. 2008, pp. 54–58.
- [12] L. Zhang, Z. Liu, and C. Honghui Xia, "Clock synchronization algorithms for network measurements," in *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1, 2002, pp. 160–169 vol.1.
- [13] J. Elson, L. Girod, and D. Estrin, "Fine-grained network time synchronization using reference broadcasts," *SIGOPS Oper. Syst. Rev.*, vol. 36, no. SI, pp. 147–163, 2002.
- [14] D. Hunt, G. Korniss, and B. Szymanski, "Network synchronization in a noisy environment with time delays: Fundamental limits and trade-offs," *Physical Review Letters*, vol. 105, no. 6, p. 068701, 2010.
- [15] H. Marouani and M. R. Dagenais, "Internal clock drift estimation in computer clusters," *J. Comp. Sys., Netw., and Comm.*, vol. 2008, pp. 9:1–9:7, Jan. 2008.
- [16] S. Moon, P. Skelly, and D. Towsley, "Estimation and removal of clock skew from network delay measurements," in *INFOCOM '99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1, Mar. 1999, pp. 227–234 vol.1.
- [17] M. Lemmon, J. Ganguly, and L. Xia, "Model-based clock synchronization in networks with drifting clocks," in *Dependable Computing, 2000. Proceedings. 2000 Pacific Rim International Symposium on*, 2000, pp. 177–184.
- [18] D. Mills, "Network time protocol (version 3) specification, implementation and analysis," 1992.
- [19] D. Xie and S. Wang, "Consensus of second-order discrete-time multi-agent systems with fixed topology," *Journal of Mathematical Analysis and Applications*, 2011.
- [20] E. Mallada and F. Paganini, "Stability of node-based multipath routing and dual congestion control," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*. IEEE, 2008, pp. 1398–1403.
- [21] W. Ren and R. W. Beard, "Consensus algorithms for double-integrator dynamics," *Distributed Consensus in Multi-vehicle Cooperative Control: Theory and Applications*, pp. 77–104, 2008.
- [22] W. Ren and R. Beard, *Distributed consensus in multi-vehicle cooperative control: theory and applications*. Springer, 2008.
- [23] S. Bhattacharyya, H. Chapellat, and L. Keel, *Robust control*. Prentice-Hall Upper Saddle River, New Jersey, 1995.
- [24] R. Horn and C. Johnson, *Matrix analysis*. Cambridge Univ Pr, 1990.

APPENDIX

A. Proof of Lemma 1

Proof: We first compute the characteristic polynomial

$$\begin{aligned} \det(\lambda I_{3n} - A) &= \begin{vmatrix} (\lambda - 1)I_n & -\tau R & \mathbf{0}_{n \times n} \\ \kappa_1 L & (\lambda - 1)I_n & \kappa_2 I_n \\ pL & 0 & (\lambda - 1 + p)I_n \end{vmatrix} \\ &= (\lambda - 1)^n \begin{vmatrix} (\lambda - 1)I_n + \frac{\tau \kappa_1}{\lambda - 1} LR & \kappa_2 I_n \\ \frac{\tau p}{\lambda - 1} LR & (\lambda - 1 + p)I_n \end{vmatrix} \\ &= \det((\lambda - 1)^2 (\lambda - 1 + p)I_n + [(\lambda - 1)\kappa_1 \\ &+ (\kappa_2 - \kappa_1)]\tau LR) = \prod_{l=1}^n g_l(\lambda), \end{aligned}$$

where $g_l(\lambda)$ is as defined in (10) and we have iteratively use the determinant property of block matrices $\det(A) =$

$\det(A_{11}) \det(A \setminus A_{11})$ where $A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}$ and $A \setminus A_{11} = A_{22} - A_{21}A_{11}^{-1}A_{12}$ is the Schur complement of A_{11} [24].

Thus, $\lambda = 1$ is a double root of the characteristic polynomial if and only if $\kappa_1 \neq \kappa_2$, $p > 0$ and τLR has a simple zero eigenvalue, i.e. (11). Now, since R is nonsingular (11) must hold for the eigenvalues of L as well, which is in fact true if and only if the directed graph $G(V, E)$ is connected [19]. ■

B. Proof of Lemma 2

Proof: We start by computing the right Jordan chain. By definition of ζ_1 , $(A - I)\zeta_1 = \mathbf{0}_n$. Thus, if $\zeta_1 = [x^T \ s^T \ y^T]^T$, then the following system of equations must be satisfied

$$\begin{aligned} \tau R s &= \mathbf{0}_n \quad (\text{a}), \quad -\kappa_1 L x - \kappa_2 y = \mathbf{0}_n \quad (\text{b}) \quad \text{and} \\ -p L x - p y &= \mathbf{0}_n \quad (\text{c}). \end{aligned} \quad (20)$$

Equation (20a) implies $s = 0$. Now, since $p > 0$, (20c) implies $Lx = -y$, which when substituted in (20b) gives $(\kappa_2 - \kappa_1)y = \mathbf{0}_n$. Thus, since $\kappa_1 \neq \kappa_2$, $y = \mathbf{0}_n$ and $x \in \ker(L)$. By choosing $x = \alpha_1 \mathbf{1}_n$ (for some $\alpha_1 \neq 0$) we obtain $\zeta_1 = \alpha_1 [\mathbf{1}_n^T \ \mathbf{0}_n^T \ \mathbf{0}_n^T]^T$.

Notice that the computation also shows that ζ_1 is the unique eigenvector of $\mu(A) = 1$ which implies that there is only one Jordan block of size 2. The second member of the chain, ζ_2 , and ζ_3 can be computed similarly by solving $(A - I_n)\zeta_2 = \zeta_1$ and $(A - (1 - p)I_n)\zeta_3 = \mathbf{0}_n$. This gives

$$\zeta_2 = \begin{bmatrix} \alpha_2 \mathbf{1}_n \\ \frac{\alpha_1}{\tau} R^{-1} \mathbf{1}_n \\ \mathbf{0}_n \end{bmatrix} \quad \text{and} \quad \zeta_3 = \alpha_3 \begin{bmatrix} -\frac{\tau \kappa_2}{p^2} \mathbf{1}_n \\ \frac{\kappa_2}{p} R^{-1} \mathbf{1}_n \\ R^{-1} \mathbf{1}_n \end{bmatrix}.$$

In computing ζ_3 , we obtain $Lx = 0$ and $Rx = -\frac{\tau}{p}s = -\frac{\kappa_2 \tau}{p^2}y$. ζ_3 follows by taking $y = \alpha_3 R^{-1} \mathbf{1}_n$.

The vectors η_1 , η_2 and η_3 can be solved in the same way using $\eta_2^T (A - I) = \mathbf{0}_n^T$, $\eta_1^T (A - I) = \eta_2^T$ and $\eta_3^T (A - (1 - p)I) = \mathbf{0}_n^T$. This gives $\eta_1 = \left[\frac{\beta_2}{\tau} R^{-1} \xi^T \ \beta_1 \xi^T \ (-\frac{\kappa_2}{p} \beta_1 + \frac{\kappa_2}{p^2} \beta_2) \xi^T \right]^T$, $\eta_2 = \beta_2 \left[\mathbf{0}_n^T \ \xi^T \ \frac{\kappa_2}{p} \xi^T \right]^T$ and $\eta_3 = \beta_3 \left[\mathbf{0}_n^T \ \mathbf{0}_n^T \ \xi^T \right]^T$. We set $\alpha_1 = \alpha_2 = \alpha_3 = 1$; this can be done WLOG provided we still satisfy $\eta_l^T \zeta_l = 1$ and $\eta_l^T \zeta_h = 0$ for $l \neq h$. Finally, $\eta_1^T \zeta_1 = 1$ gives $\beta_2 = \gamma\tau$, $\eta_3^T \zeta_3 = 1$ gives $\beta_3 = \gamma$ and $\eta_1^T \zeta_2 = 0$ gives $\beta_1 = -\beta_2 = -\gamma\tau$. ■

C. Proof of Theorem 1

Proof: We first notice that whenever $x(t_k)$ approaches (8) then

$$\lim_{h \rightarrow \infty} x(t_h) - r^* \mathbf{1}_n (t_h - t_0) = x^* \mathbf{1}_n \quad (21)$$

Sufficiency: Since we are under the assumptions of Lemmas 1 and 2 we know that $\mu(A) = 1$ has multiplicity 2 and a Jordan chain of size 2. Thus, the Jordan normal form of A is

$$A = [\zeta_1 \dots \zeta_{3n}] \begin{bmatrix} 1 & 1 & 0 & & & \\ 0 & 1 & 0 & & & \\ 0 & 0 & 1 - p & & & \\ & & & \hat{J} & & \\ & & & & \mathbf{0}_{3 \times 3(n-1)} & \\ & & & & & \hat{J} \end{bmatrix} \begin{bmatrix} \eta_1^T \\ \vdots \\ \eta_{3n}^T \end{bmatrix} \quad (22)$$

where \hat{J} has eigenvalues with spectral radius $\rho(\hat{J}) := \max_l |\mu_l(\hat{J})| < 1$. Thus, it follows that

$$\lim_{h \rightarrow \infty} A^h - \zeta_1 \eta_1^T - (h\zeta_1 + \zeta_2) \eta_2^T = \lim_{h \rightarrow \infty} [\zeta_1 \dots \zeta_{3n}] \quad (23)$$

$$\begin{bmatrix} \mathbf{0}_{2 \times 2} & \mathbf{0}_{2 \times 1} & \mathbf{0}_{2 \times (3n-2)} \\ \mathbf{0}_{1 \times 2} & (1-p)^h & \mathbf{0}_{1 \times (3n-2)} \\ \mathbf{0}_{(3n-2) \times 2} & \mathbf{0}_{(3n-2) \times 1} & \hat{J}^h \end{bmatrix} \begin{bmatrix} \eta_1^T \\ \vdots \\ \eta_{3n}^T \end{bmatrix} = \mathbf{0}_{3n}$$

where the last equality follows since $(1-p)^h \xrightarrow{h \rightarrow \infty} 0$ and

$\|\hat{J}^h\|_\varepsilon \leq \|\hat{J}\|_\varepsilon^h \leq (\rho + \varepsilon)^h \xrightarrow{h \rightarrow \infty} 0$, where the norm $\|\cdot\|_\varepsilon$ is chosen such that $\|A\|_\varepsilon = \rho(A) + \varepsilon$ [24, p. 297, Lemma 5.6.10] and ε is such $\rho(\hat{J}) + \varepsilon < 1$.

Right multiplying (23) with a given initial condition $z_0 = [x_0^T \ s_0^T \ y_0^T]^T$ and using (12) and (13) gives

$$\begin{aligned} \lim_{k \rightarrow \infty} x_k - (t_k - t_0) \gamma \mathbf{1}_n \xi^T (s_0 - \frac{\kappa_2}{p} y_0) &= \\ = \gamma \mathbf{1}_n \xi^T (R^{-1} x_0 + \tau \frac{\kappa_2}{p^2} y_0). \end{aligned} \quad (24)$$

Thus, equation (14) follows from identifying (24) and (21).

Necessity: The algorithm achieves synchronization whenever (21) holds. Then, it follows from (9) and (21) that asymptotically the system behaves according to

$$\begin{aligned} z_k &= \begin{bmatrix} x_k \\ s_k \\ y_k \end{bmatrix} = \begin{bmatrix} x^* \mathbf{1}_n \\ r^* R^{-1} \mathbf{1}_n \\ \mathbf{0}_n \end{bmatrix} + k \begin{bmatrix} \tau r^* \mathbf{1}_n \\ \mathbf{0}_n \\ \mathbf{0}_n \end{bmatrix} \\ &= (\tau r^* \zeta_2 + (x^* - \tau r^*) \zeta_1) + k r^* \tau \zeta_2. \end{aligned}$$

Thus, since P is invertible ζ_l are linearly independent. Therefore, if the system synchronizes for arbitrary initial condition, then it must be the case that the effect of the remaining modes $\mu_l(\Gamma)$ vanishes, which can only happen if for every $\mu_l(\Gamma) \neq 1$, $|\mu_l(\Gamma)| < 1$ and the multiplicity of $\mu_l(\Gamma) = 1$ is two. Now suppose that either $\kappa_1 = \kappa_2$ or $p = 0$. Then by Lemma 1, the multiplicity of $\mu_l(\Gamma) = 1$ is not two which is a contradiction. Thus, we must have $\kappa_1 \neq \kappa_2$ and $p > 0$ whenever the system synchronizes for arbitrary initial condition. ■

D. Proof of Theorem 3

Proof: We will show that when $G(V, E)$ is connected with $\mu(L) \in \mathbb{R}$, then (i)-(iii) are equivalent to the conditions of Theorem 1.

Since, $G(V, E)$ is connected and (i)-(ii) satisfies $p > 0$ and $\kappa_1 \neq \kappa_2$, the conditions of Lemma 1 are satisfied. Therefore the multiplicity of $\mu(A) = 1$ is two and by (11) these are the roots of $g_n(\lambda) = (\lambda - 1)^2(\lambda - 1 + p)$, which corresponds to the case $\nu_n = 0$.

Thus, to satisfy Theorem 1 we need to show that the remaining eigenvalues are strictly in the unit circle. This is true for the remaining root of $g_n(\lambda)$ iff (i).

For the remaining $g_l(\lambda)$, this implies that are Schur polynomials. Thus, we will show that $g_l(\lambda)$ is a Schur polynomial if and only if (i)-(iii) hold. We drop the subindex l for the rest of the proof.

We first transform the Schur stability problem into a Hurwitz stability problem. Consider the change of variable $\lambda = \frac{s+1}{s-1}$. Then $|\lambda| < 1$ if and only if $\Re[s] < 0$.

Now, since $\nu > 0$ by (11), let

$$\begin{aligned} P(s) &= \frac{(s-1)^3}{\delta \kappa p \nu} g\left(\frac{s+1}{s-1}\right) = s^3 + \left(\frac{2\kappa_1}{\delta \kappa p} - 3\right) s^2 \\ &\quad + \left(\frac{4}{\delta \kappa \nu} + 3 - \frac{4\kappa_1}{\delta \kappa p}\right) s + \frac{4(2-p)}{\delta \kappa p \nu} + \frac{2\kappa_1}{\delta \kappa p} - 1 \end{aligned}$$

where $\delta \kappa = \kappa_1 - \kappa_2$.

We will apply Hermite-Biehler Theorem to $P(s)$, but first let us express what 1) and 2) of Theorem 2 mean here.

Condition 1) becomes

$$\frac{2\kappa_1}{\delta \kappa p} - 3 > 0. \quad (25)$$

Now let $P^r(\omega)$ and $P^i(\omega)$ be as in Theorem 2, i.e. let

$$\begin{aligned} P^r(\omega) &= -\left(\frac{2\kappa_1}{\delta \kappa p} - 3\right) \omega^2 + \frac{4(2-p)}{\delta \kappa p \nu} + \frac{2\kappa_1}{\delta \kappa p} - 1 \\ P^i(\omega) &= -\omega^3 + \left(\frac{4}{\delta \kappa \nu} + 3 - \frac{4\kappa_1}{\delta \kappa p}\right) \omega \end{aligned}$$

The roots of $P^r(\omega)$ and $P^i(\omega)$ are given by $\omega_0 = \pm \sqrt{\omega_0^r}$ and $\omega_0 \in \{0, \pm \sqrt{\omega_0^i}\}$ respectively, where

$$\omega_0^r := \frac{\frac{4(2-p)}{\delta \kappa p \nu} + \frac{2\kappa_1}{\delta \kappa p} - 1}{\frac{2\kappa_1}{\delta \kappa p} - 3} \quad \text{and} \quad \omega_0^i := \frac{4}{\delta \kappa \nu} + 3 - \frac{4\kappa_1}{\delta \kappa p} \quad (26)$$

Since the roots $P^r(\omega)$ and $P^i(\omega)$ must be real, we must have $\omega_0^r > 0$ and $\omega_0^i > 0$. Therefore, by monotonicity of the square root, the interlacing condition 2) is equivalent to

$$0 < \omega_0^r < \omega_0^i. \quad (27)$$

Thus we will show: (i)-(iii) hold \iff (25) and (27) hold.

It is straightforward to see that using (i) and (ii) we can get (25). On the other hand, $\omega_0^i > 0$ from (27) together with (25) gives $0 < \frac{4}{\delta \kappa \nu} + 3 - \frac{4\kappa_1}{\delta \kappa p} < \frac{4}{\delta \kappa \nu}$, which implies that $\delta \kappa > 0$, and therefore (ii) follows.

Now using (25) and the definition of ω_0^r in (26), $\omega_0^r > 0$ becomes $\frac{4(2-p)}{\delta \kappa p \nu} + \frac{2\kappa_1}{\delta \kappa p} - 1 > 0$ which always holds under (i) and (ii) since the first term is always positive and $\frac{2\kappa_1}{\delta \kappa p} - 1 > \frac{2\kappa_1}{\delta \kappa p} - 3 > 0$ by (25).

Using (26), $\omega_0^r < \omega_0^i$ is equivalent to

$$\nu < \frac{p(\kappa_2 - \delta \kappa p)}{(\kappa_1 - \delta \kappa p)^2}. \quad (28)$$

Finally, $\nu_l = \mu_l(\tau LR) = \tau \mu_l(LR)$. Thus, since (28) should hold $\forall l \in \{1, \dots, n-1\}$, then

$$\tau < \min_l \frac{p(\kappa_2 - \delta \kappa p)}{\mu_l(LR)(\kappa_1 - \delta \kappa p)^2} = \frac{p(\kappa_2 - \delta \kappa p)}{\mu_{\max}(\kappa_1 - \delta \kappa p)^2}$$

which is exactly (iii). ■