Requirements for Secure Clock Synchronization

Lakshay Narula, and Todd Humphreys, *The University of Texas at Austin*

***Abstract*—This paper establishes a fundamental theory of secure clock synchronization. Accurate clock synchronization is the backbone of systems managing power distribution, financial transactions, telecommunication operations, database services, etc. Some clock synchronization (time transfer) systems, such as the Global Navigation Satellite Systems (GNSS), are based on one-way communication from a master to a slave clock. Others, such as the Network Transport Protocol (NTP), and the IEEE 1588 Precision Time Protocol (PTP), involve two-way communication between the master and slave. This paper shows that all one-way time transfer protocols are vulnerable to replay attacks that can potentially compromise timing information. A set of conditions for secure two-way clock synchronization is proposed and proved to be necessary and sufficient. It is shown that IEEE 1588 PTP, although a two-way synchronization protocol, is not compliant with these conditions, and is therefore insecure. Requirements for secure IEEE 1588 PTP are proposed, and a second example protocol is offered to illustrate the range of compliant systems.**

***Index Terms*—time transfer; clock synchronization; security.**

I. Introduction

Secure clock synchronization is critical to a host of technologies and infrastructure today. The phasor measurement units (PMUs) that enable monitoring and control in power grids need timing information to synchronize measurements across a wide geographical area [1]. Wireless communication networks synchronize their base stations to enable call handoff [2]. Financial networks transfer time across the globe to ensure a common time for pricing and transaction time-stamping [3]. Cloud database services such as Google's Cloud Spanner similarly require precise synchronization between the data centers to maintain consistency [4]. These clock synchronization applications have sub-millisecond accuracy and stringent security requirements.

Clock synchronization is performed either by over-the-wire packet-based communication (NTP, PTP, etc.), or by over-the-air radio signals (GNSS [2], cellular signals, LORAN [5], DCF77 [6], etc.); both wired and wireless clock synchronization are used extensively. Synchronization by GNSS is the method of choice in systems with the most stringent accuracy requirements. Equipped with atomic clocks synchronized to the most accurate time standards available, GNSS satellites can synchronize any number of stations on Earth to within a few tens of nanoseconds [7]. NTP is usually only accurate to a few milliseconds, but essentially comes for free whenever the host device is connected to a network.

One-way clock synchronization protocols are based on unidirectional communication from the time master station, `A`, to the slave station, `B`. In such protocols, `A` acts as a broadcast station and may send out timing signals either continuously or periodically. The principal drawback of one-way wireless clock synchronization protocols is their vulnerability to delay attacks in which a man-in-the-middle (MITM) adversary nefariously delays or repeats a valid transmission from one station to another. Cryptographic and other measures can improve the security of one-way protocols against delay and other signal- and data-level spoofing attacks [8]–[10], but, as will be shown, such protocols remain fundamentally insecure because of their inability to measure round trip time. They can be secured against unsophisticated attacks, but remain vulnerable to more powerful adversaries.

Two-way clock synchronization protocols involve bidirectional communication between stations `A` and `B`. Such protocols enable measurement of the round trip time of the timing signal, which is shown to be necessary for detecting MITM delay attacks. This measurement, however, is not by itself sufficient for provable security against such attacks.

This paper establishes a fundamental theory of secure clock synchronization. In contrast to the current literature on timing security [11]–[17], the problem is formalized with definitions, explicit assumptions, and proofs. The major contributions of this work are as follows:

1) One-way synchronization protocols are shown to be insecure against a MITM delay attack. Adversarial delay is shown to be indistinguishable from clock bias, and hence is unobservable without further assumptions.
2) A set of necessary conditions for secure two-way clock synchronization is presented and proved. Similar protocol-specific conditions have been previously proposed [11], [13], [18], but have not been generalized to apply to a universal clock synchronization model.
3) The proposed necessary conditions, with stricter upper bounds, are shown to be sufficient for secure synchronization in presence of a probabilistic polynomial time (PPT) adversary. Provable security for clock synchronization has not previously been explored in the literature.
4) The two-way synchronization scheme of IEEE 1588 PTP is shown to violate a necessary condition for security. This is a known vulnerability of PTP for which a fix has been proposed [11]. Having established a theory for security, this paper is able to show that the proposed fix is sufficient but is not the minimal necessary modification. A more parsimonious security requirement for PTP is presented that is both necessary and sufficient for secure synchronization.
5) A generic construction of a secure two-way clock synchronization protocol is presented to illustrate the general applicability of the proposed necessary and sufficient conditions to a range of underlying protocols.

This paper is a significant extension of [19], by the same authors: (1) the necessary conditions for security have been revamped to incorporate both continuous and packet-based clock synchronization systems, (2) a sufficiency proof for the security conditions has been formulated, and (3) protocol-specific countermeasures presented in the literature have been unified with the proposed conditions.

Wired clock synchronization is inherently more secure than its wireless counterpart because physical access to cables is easier controlled than access to radio channels. This paper primarily focuses on the more challenging task of clock synchronization over a wireless channel; nonetheless, the attacks and security protocols discussed herein also apply to wireline clock synchronization protocols in the case where the adversary gets access to the channel. For example, if an adversary is able to hijack a boundary clock in a wireline PTP network, then the resulting vulnerabilities are equivalent to that of wireless synchronization where the adversary has open access to the radio channel. In fact, an adversarial boundary clock is even more potent than a wireless adversary since it can completely block the authentic signal from reaching $\mathtt{B}$.

The rest of this paper is organized as follows. Previous works on secure clock synchronization, and their relation to this paper, are summarized in Section II. Section III presents a generic model for clock synchronization and shows that all possible one-way synchronization protocols are insecure. Section IV presents the set of security conditions for a wireless clock synchronization protocol, proving these to be necessary by contradiction. Section V presents a proof of sufficiency for the same set of conditions with stricter upper bounds. A construction of an example secure protocol is presented in Section VI, along with the security requirements for IEEE 1588 PTP. Section VII presents a simulation study of a secure clock synchronization model operating over a simplistic channel model. Concluding remarks are made in Section VIII.

II. Related Work

GNSS, NTP, and PTP are the most widely used protocols for clock synchronization. A number of research efforts have been made to assess and improve the security of these protocols. This section reviews some of the notable efforts in the literature.

The GNSS jamming and spoofing threat has been recognized in the literature for more than a decade. A survey of the current state-of-the-art in spoofing and anti-spoofing techniques is presented in [8]. Recent works on GNSS anti-spoofing techniques have specifically focused on the case of timing security. Collaborative multi-receiver [16] and direct time estimation [17] techniques have been proposed for robust GNSS clock synchronization.

The growing popularity of IEEE 1588 PTP for synchronization in critical infrastructure has brought about concerns regarding its security [11]–[15]. The threats to IEEE 1588 PTP can broadly be categorized into data-level attacks (such as modification of time messages) and physical layer attacks (such as replay and delay attacks). While cryptographic protocols are able to foil data-level attacks against realistic adversaries, some signal-level attacks, such as the delay attack, remain open vulnerabilities. Unfortunately, their execution is relatively simple. Signal-level attacks, such as the man-in-the-middle attack, have been studied in the recent past. However, these studies only include a brief discussion on countermeasure techniques, and no proof or theoretical guarantee of the efficacy of the countermeasures has been provided.

Ullman et al. [11] propose measuring the propagation delays during initialization of clock synchronization and monitoring the propagation delays during the normal operation of the time synchronization protocol. However, [11] does not prove that such a defense would be sufficient to prevent the delay attacks.

In [13], it is remarked that the clock offset computed between multiple master clocks over a symmetric channel must be zero, and thus, if multiple master clocks are available, they can detect any malicious delay introduced by an adversary. However, this defense does not consider the possibility that the adversary may only delay the packets sent to the slave nodes.

The work presented in [18] is perhaps in closest relation to the current paper. Annessi et al. upper bound the clock drift between subsequent synchronization signals using a drift model, and perform two-way exchange of timestamps such that the master clock is able to verify the time at the slave. Furthermore, given the maximum clock drift rate and the maximum and minimum propagation delay of the timing signal, they derive an upper bound on the adversarial delay that can go unnoticed. However, with conservative bounds on the maximum clock drift rate and the variation in path delays, the accuracy guarantees derived in [18] may be insufficient for certain applications. Moreover, as will be shown in this paper, they fail to take account of one the necessary conditions for secure synchronization.

This paper abstracts the clock synchronization model and assesses its security in a generic setting. It is shown that specialization of the generic security conditions to the particular protocols assessed in the aforementioned efforts leads to solutions similar or identical to those previously advanced. Thus, establishing the fundamental theory of secure clock synchronization also serves to unify the prior work in the literature.

III. System Model

A general system model for clock synchronization is shown in Fig. 1. The time seeker station, $\mathtt{B}$, wishes to synchronize its clock to that of the time master station, $\mathtt{A}$. For wireless synchronization applications, stations $\mathtt{A}$ and $\mathtt{B}$ are assumed to have known locations, $\boldsymbol{x}_{\mathtt{A}}$ and $\boldsymbol{x}_{\mathtt{B}}$, respectively. Due to clock imperfections, the time at station $\mathtt{B}$, $t_{\mathtt{B}}$, continuously drifts with respect to $t_{\mathtt{A}}$, the time at station $\mathtt{A}$. Station $\mathtt{B}$ seeks to track the relative drift of its clock by an exchange of signals between $\mathtt{A}$ and $\mathtt{B}$. Without loss of generality, this paper assumes $t_{\mathtt{A}}$ is equivalent to true time (relative to some reference epoch), a close proxy for which is GPS system time.

It is assumed that $\mathtt{A}$ and $\mathtt{B}$ are able to exchange cryptographic keys securely, if required. This exchange may occur over a public channel via a protocol such as the Diffie-Hellman key exchange [20] or via quantum key exchange techniques [21], [22]. Alternatively, symmetric keys for neighboring stations may be loaded at the time of installation.

Station $\mathtt{A}$ sends out a *sync* signal, $s_{\mathtt{A}}$, having distinct features which can be disambiguated from one another by observing a

TABLE I
NOTATION USED IN THIS PAPER

$\mathtt{A}$	Time master station
$\mathtt{B}$	Time seeker station
$t_{\mathtt{m}}^{\mathtt{m}_i}$	Transmit time, according to $\mathtt{m}$, of its ith signal feature
$t_{\mathtt{n}}^{\mathtt{m}_i}$	Receipt time, according to $\mathtt{n}$, of the ith signal feature transmitted by $\mathtt{m}$
$\tau_{\mathtt{mn}}^{i}$	Delay, in true time, experienced by the ith feature in propagating from $\mathtt{m}$ to $\mathtt{n}$
$\tau_{\mathtt{mn}_\mathcal{M}}^{i}$	Component of $\tau_{\mathtt{mn}}^{i}$ introduced by the man-in-the-middle adversary
$\tau_{\mathtt{mn}_\mathcal{N}}^{i}$	Component of $\tau_{\mathtt{mn}}^{i}$ due to natural factors, including processing, transmission, and propagation delay
$\bar{\tau}_{\mathtt{mn}}^{i}$	Modeled or *a priori* estimate of $\tau_{\mathtt{mn}_\mathcal{N}}^{i}$
$\tilde{\tau}_{\mathtt{mn}_\mathcal{N}}^{i}$	$\tau_{\mathtt{mn}_\mathcal{N}}^{i} - \bar{\tau}_{\mathtt{mn}}^{i}$
$\tau_{\mathtt{BB}}$	Delay, in true time, between the receipt of *sync* and transmission of *response* at $\mathtt{B}$
$\bar{\tau}_{\mathtt{BB}}$	Delay, according to $\mathtt{B}$, between the receipt of *sync* and transmission of *response* at $\mathtt{B}$
$\tilde{\tau}_{\mathtt{BB}}$	$\tau_{\mathtt{BB}} - \bar{\tau}_{\mathtt{BB}}$
$\Delta t_{\mathtt{AB}}^{i}$	Clock offset between $\mathtt{A}$ and $\mathtt{B}$ at the time of receipt of the ith feature at $\mathtt{B}$
$\Delta \hat{t}_{\mathtt{AB}}^{i}$	$\mathtt{B}$'s best estimate of $\Delta t_{\mathtt{AB}}^{i}$
$w_{\mathtt{mn}}^{i}$	Measurement noise associated with the measured time-of-arrival of the ith signal feature from $\mathtt{m}$ at $\mathtt{n}$
τ_{RTT}^{ij}	Round trip time, in true time, involving the ith and jth signal features of $\mathtt{A}$ and $\mathtt{B}$, respectively
$\bar{\tau}_{\mathrm{RTT}}^{ij}$	Modeled or *a priori* estimate of τ_{RTT}^{ij}
z_{RTT}^{ij}	A noisy measurement of τ_{RTT}^{ij}

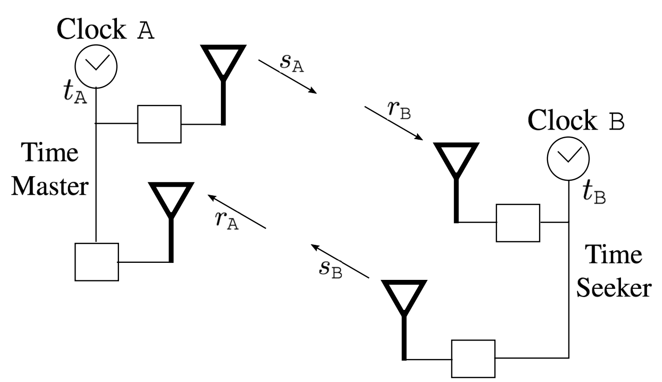

Fig. 1. Abstract model of a clock synchronization system with a time master station $\mathtt{A}$ and a time seeker station $\mathtt{B}$. The antenna outputs are driven by the clock through the receiver and transmitter blocks.

window of the signal containing the feature. The transition in $s_{\mathtt{A}}$ marking the beginning of a data packet is an example of such a signal feature. Furthermore, the system at $\mathtt{A}$ is designed such that the kth feature is transmitted at time $t_{\mathtt{A}}^{\mathtt{A}_k}$. $\mathtt{B}$ either knows $t_{\mathtt{A}}^{\mathtt{A}_k}$ by prior arrangement, or a digital representation of $t_{\mathtt{A}}^{\mathtt{A}_k}$ is encoded in $s_{\mathtt{A}}$ (e.g., a timestamp). In any case, $\mathtt{B}$ knows when the kth feature was sent, according to $\mathtt{A}$'s clock. This sets up a bijection

$$S_{\mathtt{A}}^{k} \rightleftharpoons k \rightleftharpoons t_{\mathtt{A}}^{\mathtt{A}_k} \tag{1}$$

where $S_{\mathtt{A}}^{k}$ represents a window of $s_{\mathtt{A}}$ containing the kth feature.

Station $\mathtt{B}$'s received *sync* signal, denoted $r_{\mathtt{B}}$, is a delayed and noisy replica of $s_{\mathtt{A}}$. Let $\tau_{\mathtt{AB}}^{k}$ denote the delay (in true time) experienced by the kth feature of $s_{\mathtt{A}}$ as it travels from $\mathtt{A}$ to $\mathtt{B}$. For line-of-sight (LOS) wireless communication, $\tau_{\mathtt{AB}}^{k}$ is the sum of the free-space propagation delay over the distance $\|\boldsymbol{x}_{\mathtt{B}} - \boldsymbol{x}_{\mathtt{A}}\|$ and additional delays due to interaction of the timing signal with the intervening channel.

A. One-Way Clock Synchronization Model

In one-way clock synchronization, the exchange of signals between $\mathtt{A}$ and $\mathtt{B}$ terminates with reception of the *sync* signal at $\mathtt{B}$. Let $t_{\mathtt{B}}^{\mathtt{A}_k}$ denote the time according to $\mathtt{B}$ at which the kth feature of $s_{\mathtt{A}}$ is received at $\mathtt{B}$. The window captured by $\mathtt{B}$ containing the kth feature of $s_{\mathtt{A}}$, denoted $R_{\mathtt{B}}^{k}$, enables $\mathtt{B}$ to measure $t_{\mathtt{B}}^{\mathtt{A}_k}$ to within a small error caused by measurement noise. This error, $w_{\mathtt{AB}}^{k}$, is modeled as zero-mean with variance σ_ϵ^2. The measurement itself, denoted $z_{\mathtt{B}}^{k}$, is modeled as

$$\begin{aligned} z_{\mathtt{B}}^{k} &= t_{\mathtt{B}}^{\mathtt{A}_k} + w_{\mathtt{AB}}^{k} \\ &= t_{\mathtt{A}}^{\mathtt{A}_k} + \tau_{\mathtt{AB}}^{k} - \Delta t_{\mathtt{AB}}^{k} + w_{\mathtt{AB}}^{k} \end{aligned} \tag{2}$$

where

$$\Delta t_{\mathtt{AB}}^{k} \triangleq t_{\mathtt{A}}^{\mathtt{A}_k} + \tau_{\mathtt{AB}}^{k} - t_{\mathtt{B}}^{\mathtt{A}_k} \tag{3}$$

is the unknown time offset $\mathtt{B}$ wishes to estimate. As the bijection in (1) is known to $\mathtt{B}$, $\mathtt{B}$ can obtain $t_{\mathtt{A}}^{\mathtt{A}_k}$ for the kth detected feature. If a prior estimate $\bar{\tau}_{\mathtt{AB}}^{k}$ of the delay $\tau_{\mathtt{AB}}^{k}$ is available to $\mathtt{B}$, then the desired time offset can be estimated as

$$\Delta \hat{t}_{\mathtt{AB}}^{k} = t_{\mathtt{A}}^{\mathtt{A}_k} + \bar{\tau}_{\mathtt{AB}}^{k} - z_{\mathtt{B}}^{k} \tag{4}$$

As a concrete example, consider the case of clock synchronization via GNSS in which $\mathtt{B}$ is a GNSS receiver in a known fixed location $\boldsymbol{x}_{\mathtt{B}}$, and $\mathtt{A}$ is a GNSS satellite whose location is known to vary with time as $\boldsymbol{x}_{\mathtt{A}}(t_{\mathtt{A}})$. On detection of the kth feature in a window of captured data, $\mathtt{B}$ determines $t_{\mathtt{A}}^{\mathtt{A}_k}$ using (1) and also makes the measurement

$$\begin{aligned} z_{\mathtt{B}}^{k} &= t_{\mathtt{A}}^{\mathtt{A}_k} + \tau_{\mathtt{AB}}^{k} - \Delta t_{\mathtt{AB}}^{k} + w_{\mathtt{AB}}^{k} \\ &= t_{\mathtt{A}}^{\mathtt{A}_k} + \left[\frac{\|\boldsymbol{x}_{\mathtt{B}} - \boldsymbol{x}_{\mathtt{A}}(t_{\mathtt{A}}^{\mathtt{A}_k})\| + D_\rho^k}{c}\right] - \Delta t_{\mathtt{AB}}^{k} + w_{\mathtt{AB}}^{k} \end{aligned}$$

where D_ρ^k is the sum of excess ionospheric and neutral-atmospheric delays (in distance units) and c is the speed of light.

The known receiver and satellite positions can be invoked to model the signal's propagation delay as

$$\bar{\tau}_{\mathtt{AB}}^{k} = \frac{\|\boldsymbol{x}_{\mathtt{B}} - \boldsymbol{x}_{\mathtt{A}}(t_{\mathtt{A}}^{\mathtt{A}_k})\| + \bar{D}_\rho^k}{c}$$

where $\bar{D}_\rho^k$ is a model of the excess delay D_ρ^k at the time of receipt of the kth feature at $\mathtt{B}$. The modeled excess delay is based on atmospheric models possibly refined by dual-frequency measurements [23]. An estimate of the time offset, $\Delta \hat{t}_{\mathtt{AB}}^{k}$, can then be made using $t_{\mathtt{A}}^{\mathtt{A}_k}$, $z_{\mathtt{B}}^{k}$, and $\bar{\tau}_{\mathtt{AB}}^{k}$ in (4).

It must be noted that, for one-way clock synchronization, any errors in the estimate of the distance between $\mathtt{A}$ and $\mathtt{B}$, and in the estimate of the excess channel delay, will appear as an error in the estimate of the time offset.

B. Two-Way Clock Synchronization Model

As discussed above, if an estimate of $\bar{\tau}^k_{\mathtt{AB}}$ is available, then clock synchronization is complete after $\mathtt{B}$ receives the *sync* signal $r_{\mathtt{B}}$. The *response* signal from $\mathtt{B}$ in a two-way protocol is typically used to either determine, or refine, the estimate of $\bar{\tau}^k_{\mathtt{AB}}$ with a measurement of the round trip time (RTT). The ability to measure RTT obviates the requirement that $\|\boldsymbol{x}_{\mathtt{B}} - \boldsymbol{x}_{\mathtt{A}}\|$ be known *a priori*. In IEEE 1588 PTP, for example, RTT is measured to initially obtain, and periodically refine, the value of $\bar{\tau}^k_{\mathtt{AB}}$ used in deriving $\Delta\hat{t}^k_{\mathtt{AB}}$ from (4).

In the system model considered in this paper, station $\mathtt{B}$ transmits a *response* $s_{\mathtt{B}}$ that is designed such that (1) there is a one-to-one mapping $l(k)$ between the lth feature in $s_{\mathtt{B}}$ and the kth feature in $s_{\mathtt{A}}$, and (2) the lth feature's index can be inferred by observation of a window containing it. Symbolically, if $S^l_{\mathtt{B}}$ is a window of $s_{\mathtt{B}}$ containing the lth feature of the *response* signal, then

$$S^l_{\mathtt{B}} \rightleftharpoons l(k) \rightleftharpoons k \tag{5}$$

On receipt of the kth feature in $s_{\mathtt{A}}$, at time $t^{\mathtt{A}_k}_{\mathtt{B}}$ by $\mathtt{B}$'s clock, but at $z^k_{\mathtt{B}}$ as measured by $\mathtt{B}$, $\mathtt{B}$ transmits the lth feature in $s_{\mathtt{B}}$ after a short delay, $\tau_{\mathtt{BB}}$ (in true time), hereon referred to as the *layover time*.

The layover time is introduced as a practical consideration. On receipt of $\mathtt{A}$'s kth feature, $\mathtt{B}$ is physically unable to transmit its own lth feature with zero delay. Thus, $\mathtt{B}$ is allowed to specify a short layover time, $\bar{\tau}_{\mathtt{BB}}$, after which it intends to launch its lth feature. It is important to note that the actual layover time, $\tau_{\mathtt{BB}}$, will not be the same as the intended layover time due to (1) non-zero measurement noise $w^k_{\mathtt{AB}}$ and (2) non-zero frequency offset of the clock at $\mathtt{B}$ with respect to true time. However, if the layover time is sufficiently short and the measurement noise is benign, the difference $\bar{\tau}_{\mathtt{BB}} - \tau_{\mathtt{BB}}$ can be made negligible compared to the time synchronization requirement, with the actual value depending on the quality of $\mathtt{B}$'s clock.

Station $\mathtt{A}$ receives the *response* signal as a delayed and noisy replica of $s_{\mathtt{B}}$, denoted $r_{\mathtt{A}}$. The delay experienced by the lth feature as it travels from $\mathtt{B}$ to $\mathtt{A}$, in true time, is denoted $\tau^l_{\mathtt{BA}}$. Station $\mathtt{A}$ captures a window $R^l_{\mathtt{A}}$ of $r_{\mathtt{A}}$ that enables $\mathtt{A}$ to identify the lth feature in $s_{\mathtt{B}}$ according to (5), and to infer that the received feature is in response to the kth feature transmitted by $\mathtt{A}$. Furthermore, $\mathtt{A}$ makes a noise-corrupted measurement $z^l_{\mathtt{A}}$ of the time-of-arrival of the lth feature in $s_{\mathtt{B}}$, according to $\mathtt{A}$'s clock. The noise, denoted $w^l_{\mathtt{BA}}$, is again modeled as zero-mean with variance σ^2_ϵ. The full measurement model is given by

$$\begin{aligned} z^l_{\mathtt{A}} &= t^{\mathtt{B}_l}_{\mathtt{A}} + w^l_{\mathtt{BA}} \\ &= t^{\mathtt{A}_k}_{\mathtt{A}} + \tau^k_{\mathtt{AB}} + \tau_{\mathtt{BB}} + \tau^l_{\mathtt{BA}} + w^l_{\mathtt{BA}} \end{aligned}$$

Since $t^{\mathtt{A}_k}_{\mathtt{A}}$ is exactly known at $\mathtt{A}$, a direct noisy measurement of the round trip time $\tau^k_{\mathtt{AB}} + \tau_{\mathtt{BB}} + \tau^l_{\mathtt{BA}}$ can be made as

$$z^{kl}_{\mathrm{RTT}} \triangleq z^l_{\mathtt{A}} - t^{\mathtt{A}_k}_{\mathtt{A}} \tag{6}$$

Note that the noise $w^l_{\mathtt{BA}}$ and $w^k_{\mathtt{AB}}$ in z^{kl}_{RTT} is embedded within $z^l_{\mathtt{A}}$ and $\tau_{\mathtt{BB}}$, respectively. Under the assumption of symmetric delays, i.e., $\tau^k_{\mathtt{AB}} = \tau^l_{\mathtt{BA}}$, and with knowledge of $\bar{\tau}_{\mathtt{BB}}$, the measured RTT in (6) can be exploited to improve the modeled propagation delay for future exchanges between $\mathtt{A}$ and $\mathtt{B}$:

$$\bar{\tau}^m_{\mathtt{AB}} = \bar{\tau}^n_{\mathtt{BA}} = \frac{z^{kl}_{\mathrm{RTT}} - \bar{\tau}_{\mathtt{BB}}}{2}$$

where $m > k$ and $n > l$.

The two-way exchange of *sync* and *response* messages is summarized visually in Fig. 2.

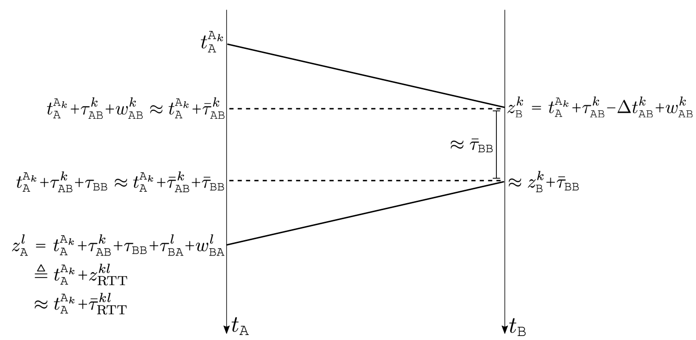

Fig. 2. Two-way exchange of *sync* and *response* messages between $\mathtt{A}$ and $\mathtt{B}$ in the absence of a man-in-the-middle adversary.

Since RTT will play a central role in the discussion on secure synchronization later on, various definitions and assumptions concerning RTT are stated here for clarity:

- RTT for the kth feature in $s_{\mathtt{A}}$ and the corresponding lth feature in $s_{\mathtt{B}}$ is defined as

  $$\tau^{kl}_{\mathrm{RTT}} \triangleq \tau^k_{\mathtt{AB}} + \tau_{\mathtt{BB}} + \tau^l_{\mathtt{BA}}$$

- Measured RTT includes, in addition to RTT, measurement noise at $\mathtt{A}$; it is modeled as

  $$z^{kl}_{\mathrm{RTT}} = \tau^k_{\mathtt{AB}} + \tau_{\mathtt{BB}} + \tau^l_{\mathtt{BA}} + w^l_{\mathtt{BA}}$$

- Modeled RTT, also called the prior estimate of RTT, is defined as

  $$\bar{\tau}^{kl}_{\mathrm{RTT}} \triangleq \bar{\tau}^k_{\mathtt{AB}} + \bar{\tau}_{\mathtt{BB}} + \bar{\tau}^l_{\mathtt{BA}} \tag{7}$$

  For example, in the case of wireless clock synchronization with LOS electromagnetic signals, a prior estimate of RTT is based on the distance between $\mathtt{A}$ and $\mathtt{B}$ and on models of channel delays in excess of free-space propagation between these.
- The modeled RTT, $\bar{\tau}^{kl}_{\mathrm{RTT}}$, can be refined with measurements of RTT in a two-way protocol. Alternatively, as will be discussed later, if an accurate modeled RTT is available, it and the measured RTT can be used to detect delay attacks.
- Unambiguous measurement of RTT requires that there exist a one-to-one mapping between the signal features in $s_{\mathtt{A}}$ and $s_{\mathtt{B}}$, as mathematically represented in (5). On detection of the lth feature in $s_{\mathtt{B}}$, $\mathtt{A}$ must be able to deduce that this feature was transmitted approximately $\bar{\tau}_{\mathtt{BB}}$ after $\mathtt{B}$ received the kth feature in $s_{\mathtt{A}}$. This requirement is appropriately a part of the RTT definition since it enables $\mathtt{A}$ to unambiguously measure RTT.

C. Attack Model

The attack model in this paper considers a MITM adversary $\mathcal{M}$. The available computational resources allow $\mathcal{M}$ to execute probabilistic polynomial time (PPT) algorithms. $\mathcal{M}$ can receive, detect, and replay signals from $\mathtt{A}$ and $\mathtt{B}$ with arbitrarily precise directional antennas. Additionally, $\mathcal{M}$ has precise knowledge of $\boldsymbol{x}_{\mathtt{A}}$ and $\boldsymbol{x}_{\mathtt{B}}$, and can take up any position around or between the two stations. It has unrestricted access to the signals that $\mathtt{A}$ and $\mathtt{B}$ exchange over the air, and has complete knowledge of their synchronization protocol save for the cryptographic keys.

Let L denote the alert limit, defined as the error in time synchronization not to be exceeded without issuing an alert.

Definition III.1. *Clock synchronization is defined to be compromised if* $|\Delta t_{\mathtt{AB}} - \Delta \hat{t}_{\mathtt{AB}}| \geq L$.

Note that, in the absence of an adversary, clock synchronization is not compromised so long as

$$|\tau^k_{\mathtt{AB}} - \bar{\tau}^k_{\mathtt{AB}} + w^k_{\mathtt{AB}}| < L$$

However, in the presence of a MITM adversary, the *sync* signal is delayed or advanced such that

$$\tau^k_{\mathtt{AB}} = \tau^k_{\mathtt{AB}_{\mathcal{N}}} + \tau^k_{\mathtt{AB}_{\mathcal{M}}} \tag{8}$$

where $\tau^k_{\mathtt{AB}_{\mathcal{N}}} > 0$ is the natural or physical delay (equal to $\tau^k_{\mathtt{AB}}$ in the absence of an adversary) and $\tau^k_{\mathtt{AB}_{\mathcal{M}}} \geq 0$ is the adversarial delay. In this case, if

$$|\tau^k_{\mathtt{AB}} - \bar{\tau}^k_{\mathtt{AB}} + w^k_{\mathtt{AB}}| = |\tau^k_{\mathtt{AB}_{\mathcal{N}}} - \bar{\tau}^k_{\mathtt{AB}} + \tau^k_{\mathtt{AB}_{\mathcal{M}}} + w^k_{\mathtt{AB}}| \geq L \tag{9}$$

then clock synchronization is compromised.

D. Vulnerability of One-Way Clock Synchronization

One-way clock synchronization is fundamentally vulnerable to a delay attack because it provides no mechanism to measure RTT. The adversary $\mathcal{M}$ can compromise any one-way wireless clock synchronization protocol by retransmitting the authentic *sync* signal from $\mathtt{A}$ such that the retransmitted signal, $s_{\mathcal{M}}$, overpowers or otherwise supersedes the authentic signal $s_{\mathtt{A}}$. In the absence of additional assumptions beyond those underpinning the one-way protocol described earlier, $\mathcal{M}$ can introduce an arbitrary delay $\tau^k_{\mathtt{AB}_{\mathcal{M}}}$ in its retransmission, thereby compromising the synchronization process.

Note that whereas counterfeit signal attacks can be prevented by authentication and cryptographic methods [24], these techniques do not prevent delay attacks because the delayed or repeated signal has the same cryptographic characteristics as that of the genuine signal, the only difference being that it is received with a (possibly small) additional delay.

The delay introduced by $\mathcal{M}$ is added to the natural delay, $\tau^k_{\mathtt{AB}_{\mathcal{N}}}$, of the signal between $\mathtt{A}$ and $\mathtt{B}$. As a result, an error of $\approx \tau^k_{\mathtt{AB}_{\mathcal{M}}}$ is introduced in the estimated time offset at $\mathtt{B}$. From (4), it follows that

$$\begin{aligned}\Delta \hat{t}^k_{\mathtt{AB}} &= t^{\mathtt{A}_k}_{\mathtt{A}} + \bar{\tau}^k_{\mathtt{AB}} - z^k_{\mathtt{B}} \\ &= t^{\mathtt{A}_k}_{\mathtt{A}} + \bar{\tau}^k_{\mathtt{AB}} - (t^{\mathtt{A}_k}_{\mathtt{A}} + \tau^k_{\mathtt{AB}} - \Delta t^k_{\mathtt{AB}} + w^k_{\mathtt{AB}}) \\ &= (\bar{\tau}^k_{\mathtt{AB}} - \tau^k_{\mathtt{AB}_{\mathcal{N}}}) - \tau^k_{\mathtt{AB}_{\mathcal{M}}} + \Delta t^k_{\mathtt{AB}} - w^k_{\mathtt{AB}} \\ &\approx \Delta t^k_{\mathtt{AB}} - \tau^k_{\mathtt{AB}_{\mathcal{M}}}\end{aligned} \tag{10}$$

where it is assumed that the error due to inaccurately modeled delay is negligible and that $\sigma_\epsilon \ll \tau^k_{\mathtt{AB}_{\mathcal{M}}}$. In the absence of an RTT measurement, and without further assumptions on the nature of the protocol or the clock drift model considered, the adversarial delay $\tau_{\mathtt{AB}_{\mathcal{M}}}$ is indistinguishable from a clock offset of the same magnitude.

To be sure, measures can be taken to make a MITM delay attack harder to execute without detection. But, importantly, these measures cannot guarantee that the synchronization will remain uncompromised. Various measures proposed in the literature, and their shortcomings, are discussed below.

a) Received Signal Strength Monitoring: The adversary $\mathcal{M}$ might attempt to overpower the authentic signal in order to spoof the *sync* message, leading to an increase in the total signal power received at $\mathtt{B}$. Station $\mathtt{B}$ could monitor the received signal strength (RSS) to detect such an attack [25]. However, a potent adversary could transmit, in addition to its delayed signal, an amplitude-matched, phase-inverted nulling signal that annihilates the authentic *sync* signal $s_{\mathtt{A}}$ as received at $\mathtt{B}$, thus preventing an unusual increase in received power at $\mathtt{B}$. If $\mathcal{M}$ is positioned along the straight-line path between $\mathtt{A}$ and $\mathtt{B}$, nulling of $s_{\mathtt{A}}$ can be effected without prior knowledge of $s_{\mathtt{A}}$. A laboratory demonstration of such nulling is reported in [26].

b) Selective Rejection of False Signal: If $\mathtt{B}$ receives both the authentic and false (delayed) *sync* signals, it may be able to apply angle-of-arrival or signal processing techniques to selectively reject the delayed signal [8], [9], [27], [28]. However, discrimination based on angle-of-arrival fails if $\mathcal{M}$ is positioned along the line from $\mathtt{A}$ to $\mathtt{B}$, and, as conceded in [9], signal-processing-based techniques for selective rejection of false signals can be thwarted by an adversary transmitting an additional nulling signal, as described above.

c) Collaborative Verification: Multiple time seekers may attempt to synchronize to the same time master. In this scenario, the time seekers can potentially detect malicious activity by cross-checking the received signals [16]. In the simplest implementation, all time seekers can collaborate to verify that they are synchronized amongst each other. In case of an uncoordinated attack against a subset of time seekers, this verification would expose the attack since the time offset computed at the attacked subset would be different from that computed at the other stations. In principle, however, it is possible for an adversary to execute a coordinated attack against all the time seekers, thus concealing its presence.

IV. Necessary Conditions for Secure Synchronization

This section presents a set of conditions for secure two-way clock synchronization and proves these to be necessary by contradiction. In other words, it is shown that if a two-way clock synchronization protocol does not satisfy any one of these proposed conditions, there exists an attack that can compromise clock synchronization without detection.

It is important to note that the ability to measure RTT in a two-way protocol is necessary, but not sufficient, for provably secure synchronization. As an example, IEEE 1588 PTP is a two-way protocol that has been proposed as an alternative to GNSS for sub-microsecond clock synchronization in critical infrastructure such as the PMU network. But, despite the bi-directional exchange between stations, and hence the ability to measure RTT, recent work has shown that PTP is vulnerable to delay attacks in which a MITM introduces asymmetric delay between $\mathtt{A}$ and $\mathtt{B}$. Asymmetric delay breaks the assumption that $\tau^k_{\mathtt{AB}} = \tau^k_{\mathtt{BA}}$ and leads to an erroneous prior for $\bar{\tau}_{\mathtt{AB}}$ and $\bar{\tau}_{\mathtt{BA}}$ for future exchanges. This vulnerability is documented in both the literature [11], [13], [18] and the IEEE 1588-2008 standard. Thus, a secure two-way clock synchronization protocol must satisfy additional security requirements beyond the ability to measure RTT.

The conditions introduced below are not tied to any specific protocol, unlike some measures proposed in the current literature [11]–[17]. They are generally applicable to any two-way protocol (e.g., PTP) for which the foregoing two-way synchronization model applies.

Assuming the time master $\mathtt{A}$ initiates the two-way communication, the necessary conditions for secure clock synchronization are as follows:

1) Both $\mathtt{A}$ and $\mathtt{B}$ must transmit unpredictable waveforms to prevent the adversary $\mathcal{M}$ from generating counterfeit signals that pass authentication. In practice, this implies the use of a cryptographic construct such as a message authentication code (MAC) or a digital signature.
2) The propagation time of the signal must be irreducible to within the alert limit L along both signal paths. For wireless clock synchronization, this condition implies synchronization via LOS electromagnetic signals as $L \to 0$.
3) The RTT between $\mathtt{A}$ and $\mathtt{B}$ must be known to $\mathtt{A}$ and measurable by $\mathtt{A}$ to within the alert limit L. The RTT must include the delays internal to both $\mathtt{A}$ and $\mathtt{B}$, in addition to the propagation delay. Station $\mathtt{A}$ must know of any intentional delay introduced by $\mathtt{B}$, such as the layover time $\tau_{\mathtt{BB}}$ introduced earlier.

A. *Proof of Necessity of Conditions*

1) Stations $\mathtt{A}$ and $\mathtt{B}$ must transmit unpredictable signals: To prove this condition is necessary, two scenarios are considered: *a)* station $\mathtt{A}$ transmits a signal waveform $s_{\mathtt{A}}$ that is predictable, and, *b)* station $\mathtt{B}$ transmits a signal waveform $s_{\mathtt{B}}$ that is predictable.

a) $s_{\mathtt{A}}$ is predictable: $\mathcal{M}$ can compromise synchronization without detection as follows:

i) $\mathcal{M}$ takes up a position between $\mathtt{A}$ and $\mathtt{B}$ along the line joining the antennas at the two stations.
ii) $\mathcal{M}$ initially transmits a replica of $s_{\mathtt{A}}$ such that $\mathtt{B}$ receives identical signals from both $\mathtt{A}$ and $\mathcal{M}$. Subsequently, $\mathcal{M}$ increases its signal power or otherwise supersedes $s_{\mathtt{A}}$ (e.g., via signal nulling, as discussed earlier) such that $\mathtt{B}$ tracks $s_{\mathcal{M}}$, the signal transmitted by $\mathcal{M}$. (Hereafter, whenever signals from $\mathcal{M}$ compete with those from $\mathtt{A}$ or $\mathtt{B}$, it will be assumed that those from $\mathcal{M}$ exert control.)
iii) Exploiting the predictability of $s_{\mathtt{A}}$, $\mathcal{M}$ advances its replica $s_{\mathcal{M}}$ with respect to $s_{\mathtt{A}}$ by $|\tau^k_{\mathtt{AB}_{\mathcal{M}}}|$, where $\tau^k_{\mathtt{AB}_{\mathcal{M}}} < 0$. $\mathtt{B}$ tracks the advanced signal, resulting in an error of $\tau^k_{\mathtt{AB}_{\mathcal{M}}}$ in the computed $\Delta\hat{t}^k_{\mathtt{AB}}$ as shown in (10).
iv) $\mathtt{B}$ transmits the unpredictable *response* $s_{\mathtt{B}}$ compliant with the pre-arranged layover time $\bar{\tau}_{\mathtt{BB}}$. $\mathcal{M}$ intercepts this signal from $\mathtt{B}$, and replays it to $\mathtt{A}$ with a delay of $\tau^l_{\mathtt{BA}_{\mathcal{M}}} = -\tau^k_{\mathtt{AB}_{\mathcal{M}}} > 0$, causing $\mathtt{A}$ to track the delayed signal. As a result, the RTT is $\tau^k_{\mathtt{AB}} + \tau_{\mathtt{BB}} + \tau^l_{\mathtt{BA}}$ as $\mathtt{A}$ expects. In summary:

$$\begin{aligned} \tau^k_{\mathtt{AB}} &= \tau^k_{\mathtt{AB}_{\mathcal{N}}} + \tau^k_{\mathtt{AB}_{\mathcal{M}}} \\ \tau^l_{\mathtt{BA}} &= \tau^l_{\mathtt{BA}_{\mathcal{N}}} + \tau^l_{\mathtt{BA}_{\mathcal{M}}} = \tau^l_{\mathtt{BA}_{\mathcal{N}}} - \tau^k_{\mathtt{AB}_{\mathcal{M}}} \\ \Rightarrow \tau^k_{\mathtt{AB}} + \tau^l_{\mathtt{BA}} &= \tau^k_{\mathtt{AB}_{\mathcal{N}}} + \tau^l_{\mathtt{BA}_{\mathcal{N}}} \end{aligned}$$

Thus, $\mathcal{M}$ undoes the effect of its *sync* advance, preventing $\mathtt{A}$ from detecting the attack.

b) $s_{\mathtt{B}}$ is predictable: $\mathcal{M}$ can compromise synchronization without detection by replicating $\mathtt{B}$'s behavior:

i) $\mathcal{M}$ takes up a position between $\mathtt{A}$ and $\mathtt{B}$ along the line joining the antennas at the two stations.
ii) $\mathcal{M}$ receives the *sync* signal and generates a valid *response* with a delay

$$\bar{\tau}_{\mathtt{BB}} + \frac{\|\boldsymbol{x}_{\mathcal{M}} - \boldsymbol{x}_{\mathtt{B}}\|}{\|\boldsymbol{x}_{\mathtt{A}} - \boldsymbol{x}_{\mathtt{B}}\|}\left(\bar{\tau}^k_{\mathtt{AB}} + \bar{\tau}^l_{\mathtt{BA}}\right) \tag{11}$$

such that the RTT is $\bar{\tau}^k_{\mathtt{AB}} + \bar{\tau}_{\mathtt{BB}} + \bar{\tau}^l_{\mathtt{BA}}$, as $\mathtt{A}$ expects.
iii) $\mathcal{M}$ records the unpredictable signal from $\mathtt{A}$ and replays it to $\mathtt{B}$ with an arbitrary delay $\tau^k_{\mathtt{AB}_{\mathcal{M}}} > 0$. This results in an error of approximately $\tau^k_{\mathtt{AB}_{\mathcal{M}}}$ in the computed $\Delta\hat{t}^k_{\mathtt{AB}}$ at $\mathtt{B}$, as shown in (10).

2) Propagation time must be irreducible to within L: If there exists a channel that reduces the propagation time from $\mathtt{A}$ to $\mathtt{B}$ or from $\mathtt{B}$ to $\mathtt{A}$ by more than L as compared to the channel used by $\mathtt{A}$ and $\mathtt{B}$, then $\mathcal{M}$ can compromise synchronization without detection. The following attack assumes the propagation time from $\mathtt{A}$ to $\mathtt{B}$ is reducible by more than L; a similar attack exploits the situation in which the propagation time from $\mathtt{B}$ to $\mathtt{A}$ is reducible by more than L.

i) $\mathcal{M}$ records the *sync* signal $s_{\mathtt{A}}$ going from $\mathtt{A}$ to $\mathtt{B}$.
ii) $\mathcal{M}$ makes the recorded signal reach $\mathtt{B}$ advanced by $|\tau^k_{\mathtt{AB}_{\mathcal{M}}}|$ compared to $s_{\mathtt{A}}$, where $\tau^k_{\mathtt{AB}_{\mathcal{M}}} < -L$. An error of $\tau^k_{\mathtt{AB}_{\mathcal{M}}}$ is introduced in the time offset value computed at $\mathtt{B}$ as shown in (10).
iii) $\mathcal{M}$ records the *response* signal $s_{\mathtt{B}}$, which has the expected pre-arranged layover time $\tau_{\mathtt{BB}} \approx \bar{\tau}_{\mathtt{BB}}$. $\mathcal{M}$ replays this signal to $\mathtt{A}$ with a delay of $\tau^l_{\mathtt{BA}_{\mathcal{M}}} = -\tau^k_{\mathtt{AB}_{\mathcal{M}}}$ such that the RTT is consistent with what $\mathtt{A}$ expects.

3) RTT known to and measurable by $\mathtt{A}$ to within L: Synchronization can be compromised without detection if $|z^{kl}_{\mathrm{RTT}} - \bar{\tau}^{kl}_{\mathrm{RTT}}| > L$ with non-negligible probability even in the absence of an adversary. This condition can be met if *a)* the prior estimates $\bar{\tau}^k_{\mathtt{AB}}$, $\bar{\tau}^l_{\mathtt{BA}}$, or $\bar{\tau}_{\mathtt{BB}}$ are not accurate to the corresponding true values to within L, or *b)* the magnitude of the measurement error sum $|w^k_{\mathtt{AB}} + w^l_{\mathtt{BA}}|$ is larger than L. Note that the condition $|w^k_{\mathtt{AB}}| > L$ compromises synchronization even absent an adversary. An adversary $\mathcal{M}$ can exploit the condition $|z^{kl}_{\mathrm{RTT}} - \bar{\tau}^{kl}_{\mathrm{RTT}}| > L$ as follows:

i) $\mathcal{M}$ initially transmits a replica of $s_{\mathtt{A}}$ such that $\mathtt{B}$ receives identical signals from both $\mathtt{A}$ and $\mathcal{M}$. Subsequently, $\mathcal{M}$

introduces a delay $\tau^k_{\mathrm{AB}_\mathcal{M}} > 0$ in the replayed signal $s_\mathcal{M}$. As assumed earlier, $s_\mathcal{M}$ exerts control and introduces an error of approximately $\tau^k_{\mathrm{AB}_\mathcal{M}}$ in the computed $\Delta\hat{t}^k_{\mathrm{AB}}$ at B, as shown in (10).

ii) Station B transmits the *response* signal with the pre-arranged layover time $\tau_{\mathrm{BB}} \approx \bar{\tau}_{\mathrm{BB}}$ with respect to the delayed signal.

iii) In the received signal r_{A}, A identifies the expected feature $l(k)$. The RTT, if measurable, includes the delay $\tau^k_{\mathrm{AB}_\mathcal{M}}$ introduced by $\mathcal{M}$.

iv) However, A is unable to definitively declare an attack, since the errors in the modeled RTT and/or the measurement of RTT are possibly larger than L. In other words, it is not possible to claim that $|z^{kl}_{\mathrm{RTT}} - \bar{\tau}^{kl}_{\mathrm{RTT}}| > L$ only in the presence of adversarial delay.

V. Proof of Sufficiency

This section presents a sufficiency proof for the set of security conditions proposed in the previous section. A sufficiency proof guarantees secure synchronization under the considered system and attack models. This paper draws inspiration from the literature on modern cryptography and formalizes the problem of secure clock synchronization with explicit definitions, assumptions, and proofs.

A. Assumptions

This proof assumes that the system under consideration strictly complies with the set of necessary security conditions. Specifically,

1) Both A and B use an authenticated encryption scheme to generate unpredictable and verifiably authentic signals in the presence of a probabilistic polynomial time (PPT) adversary.
2) The difference between the RTT along the communication channel between A and B and the shortest possible RTT is negligible as compared to L.
3) The difference between the modeled delays $\bar{\tau}^k_{\mathrm{AB}}$ and $\bar{\tau}^l_{\mathrm{BA}}$ and the true delays τ^k_{AB} and τ^l_{BA}, respectively, is negligible as compared to L.

$$|\bar{\tau}^k_{\mathrm{AB}} - \tau^k_{\mathrm{AB}_\mathcal{N}}| \ll L \tag{12}$$

and

$$|\bar{\tau}^l_{\mathrm{BA}} - \tau^l_{\mathrm{BA}_\mathcal{N}}| \ll L \tag{13}$$

Furthermore, A and B agree upon a fixed layover time $\bar{\tau}_{\mathrm{BB}}$, and the difference between this and the true layover time is negligible: $|\tau_{\mathrm{BB}} - \bar{\tau}_{\mathrm{BB}}| \ll L$.

4) The standard deviation of the noise corrupting the measurements $t^{\mathrm{A}_k}_{\mathrm{B}}$ and $t^{\mathrm{B}_l}_{\mathrm{A}}$ is negligible compared to the alert limit:

$$\sigma_\epsilon \ll L \tag{14}$$

Notice that the above assumptions are the same as the necessary conditions in Section IV, but with stricter upper bounds on the conditions.

If symmetric keys are exchanged prior to synchronization, then private-key cryptographic schemes such as Encrypt-then-MAC [29] can be used for authenticated encryption. Alternatively, if the keys must be exchanged over a public channel, then digital signatures [30] can be used to authenticate the encrypted messages. Cryptographic authentication schemes like MAC and digital signatures generate a tag associated with a message. Qualitatively, a MAC or digital signature scheme is secure if a PPT adversary, even when given access to multiple valid message-tag pairs of its own choice (as many as possible in polynomial time), cannot generate a valid tag for a new message with non-negligible probability. Irrespective of the cryptographic scheme used, this proof assumes that the probability of $\mathcal{M}$ generating a new valid *sync* or *response* signal is a negligible function of the key length n:

$$\mathbb{P}\left[\mathsf{Valid}\right] < \mathsf{negl}(n) \tag{15}$$

To detect an attack before the synchronization error exceeds L, A must select a threshold lower than L beyond which an attack is declared. Consider the modeled RTT, $\bar{\tau}^{kl}_{\mathrm{RTT}}$, as defined in (7), and the measurement z^{kl}_{RTT} as defined in (6). A threshold less than L, say $L - \delta$ with $0 < \delta < L$, is set by station A such that if $|z^{kl}_{\mathrm{RTT}} - \bar{\tau}^{kl}_{\mathrm{RTT}}| > L - \delta$, then an attack is declared.

B. Definitions

Definition V.1. *A PPT adversary $\mathcal{M}$ succeeds if clock synchronization is compromised (Definition III.1) and*

$$|z^{kl}_{\mathrm{RTT}} - \bar{\tau}^{kl}_{\mathrm{RTT}}| \le L - \delta$$

Definition V.2. *Faster-than-light (superluminal) propagation is defined to be hard if $\mathcal{M}$ cannot propagate a signal at a speed higher than the speed of light with non-negligible probability. Under hardness of superluminal propagation*

$$\mathbb{P}[\mathsf{Superluminal}] \approx 0$$

Definition V.3. *A clock synchronization protocol is defined to be secure if, under the hardness of superluminal propagation assumption,*

$$\mathbb{P}[\mathsf{Success}] < \mathsf{negl}(n)$$

where Success *for $\mathcal{M}$ is defined in Definition V.1.*

C. Proof

In the presence of an adversary $\mathcal{M}$, the measurement z^{kl}_{RTT} is modeled as

$$z^{kl}_{\mathrm{RTT}} = \tau^k_{\mathrm{AB}_\mathcal{N}} + \tau^k_{\mathrm{AB}_\mathcal{M}} + \tau^l_{\mathrm{BA}_\mathcal{N}} + \tau^l_{\mathrm{BA}_\mathcal{M}} + \tau_{\mathrm{BB}} + w^l_{\mathrm{BA}} \tag{16}$$

Let $\tilde{\tau}^k_{\mathrm{AB}_\mathcal{N}}$ and $\tilde{\tau}^l_{\mathrm{BA}_\mathcal{N}}$ denote the error in the modeled signal delay due to natural/physical phenomenon. Also, let $\tilde{\tau}_{\mathrm{B}}$ be the difference between the intended layover time $\bar{\tau}_{\mathrm{BB}}$ and the actual layover time τ_{BB}. Note that these might be positive or negative.

$$\tilde{\tau}^k_{\mathrm{AB}_\mathcal{N}} = \tau^k_{\mathrm{AB}_\mathcal{N}} - \bar{\tau}^k_{\mathrm{AB}} \tag{17}$$

$$\tilde{\tau}^l_{\mathrm{BA}_\mathcal{N}} = \tau^l_{\mathrm{BA}_\mathcal{N}} - \bar{\tau}^l_{\mathrm{BA}} \tag{18}$$

$$\tilde{\tau}_{\mathrm{BB}} = \tau_{\mathrm{BB}} - \bar{\tau}_{\mathrm{BB}} \tag{19}$$

From (7), (16), (17), (18), and (19) it follows that

$$z^{kl}_{\mathrm{RTT}} = \bar{\tau}^{kl}_{\mathrm{RTT}} + \tilde{\tau}^{k}_{\mathrm{AB}_{\mathcal{N}}} + \tau^{k}_{\mathrm{AB}_{\mathcal{M}}} + \tilde{\tau}^{l}_{\mathrm{BA}_{\mathcal{N}}} + \tau^{l}_{\mathrm{BA}_{\mathcal{M}}} + \tilde{\tau}_{\mathrm{BB}} + w^{l}_{\mathrm{BA}}$$

Following the assumptions in (12) and (13), the residual delays are negligible in comparison to L:

$$|\tilde{\tau}^{k}_{\mathrm{AB}_{\mathcal{N}}}| \ll L \tag{20}$$

$$|\tilde{\tau}^{l}_{\mathrm{BA}_{\mathcal{N}}}| \ll L \tag{21}$$

This assumption is reasonable since otherwise the system could not confidently meet the accuracy requirements even in the absence of an adversary. Also, if $\bar{\tau}_{\mathrm{BB}}$ is a short time interval and the measurement noise σ_{ϵ} is benign, it is reasonable to assume that

$$|\tilde{\tau}_{\mathrm{BB}}| \ll L \tag{22}$$

Note that $\mathcal{M}$ can advance the signal by (*a*) forging a valid message/tag pair, or (*b*) propagating the signal faster-than-light. The assumptions of secure MAC and hardness of superluminal propagation enforce that

$$\begin{aligned}\mathbb{P}[\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} < 0] &< \mathbb{P}[\mathsf{Valid}] + \mathbb{P}[\mathsf{Superluminal}] \\ &\approx \mathsf{negl}(n)\end{aligned}$$

In order to stay undetected, the adversary must ensure

$$\begin{aligned}L - \delta &\geq |z^{kl}_{\mathrm{RTT}} - \bar{\tau}^{kl}_{\mathrm{RTT}}| \\ &= |\tilde{\tau}^{k}_{\mathrm{AB}_{\mathcal{N}}} + \tau^{k}_{\mathrm{AB}_{\mathcal{M}}} + \tilde{\tau}^{l}_{\mathrm{BA}_{\mathcal{N}}} + \tau^{l}_{\mathrm{BA}_{\mathcal{M}}} + \tilde{\tau}_{\mathrm{BB}} + w^{l}_{\mathrm{BA}}| \end{aligned} \tag{23}$$

At the same time, in order to compromise time transfer, from (9), $\mathcal{M}$ must ensure

$$\begin{aligned} L &\leq |\tilde{\tau}^{k}_{\mathrm{AB}_{\mathcal{N}}} + \tau^{k}_{\mathrm{AB}_{\mathcal{M}}} + w^{k}_{\mathrm{AB}}| \\ &\leq |\tilde{\tau}^{k}_{\mathrm{AB}_{\mathcal{N}}} + w^{k}_{\mathrm{AB}}| + |\tau^{k}_{\mathrm{AB}_{\mathcal{M}}}| \\ \Rightarrow |\tau^{k}_{\mathrm{AB}_{\mathcal{M}}}| &\geq L - |\tilde{\tau}^{k}_{\mathrm{AB}_{\mathcal{N}}} + w^{k}_{\mathrm{AB}}| \end{aligned} \tag{24}$$

The probability of success for $\mathcal{M}$ is evaluated as

$$\begin{aligned}\mathbb{P}[\mathsf{Success}] &= \mathbb{P}[(\mathsf{Success}) \cap (\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} < 0)] + \\ &\quad \mathbb{P}[(\mathsf{Success}) \cap (\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} \geq 0)] \\ &= \mathbb{P}[(\mathsf{Success}) | (\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} < 0)]\mathbb{P}[\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} < 0] + \\ &\quad \mathbb{P}[(\mathsf{Success}) \cap (\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} \geq 0)] \\ &\leq \mathbb{P}[\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} < 0] + \mathbb{P}[(\mathsf{Success}) \cap (\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} \geq 0)] \\ &< \mathsf{negl}(n) + \mathbb{P}[(\mathsf{Success}) \cap (\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} \geq 0)] \end{aligned} \tag{25}$$

In the case where $\tau_{\mathrm{AB}_{\mathcal{M}}} \geq 0$, (24) simplifies to

$$\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} \geq L - |\tilde{\tau}^{k}_{\mathrm{AB}_{\mathcal{N}}} + w^{k}_{\mathrm{AB}}|$$

Substituting the least possible value of $\tau^{k}_{\mathrm{AB}_{\mathcal{M}}}$ into (23), it follows that

$$|\tilde{\tau}^{k}_{\mathrm{AB}_{\mathcal{N}}} + L - |\tilde{\tau}^{k}_{\mathrm{AB}_{\mathcal{N}}} + w^{k}_{\mathrm{AB}}| + \tilde{\tau}^{l}_{\mathrm{BA}_{\mathcal{N}}} + \tau^{l}_{\mathrm{BA}_{\mathcal{M}}} + \tilde{\tau}_{\mathrm{BB}} + w^{l}_{\mathrm{BA}}| \leq L - \delta$$

Notice that from the assumptions made in (14), (20), (21), and (22), all terms except L and $\tau^{l}_{\mathrm{BA}_{\mathcal{M}}}$ on the left-hand side of the inequality are negligible compared to L; thus,

$$|L + \tau^{l}_{\mathrm{BA}_{\mathcal{M}}}| \leq L - \delta$$

Since both L and $L - \delta$ are defined to be positive, the above inequality simplifies to

$$\tau^{l}_{\mathrm{BA}_{\mathcal{M}}} \leq -\delta$$

where $\delta > 0$. Thus, for $\mathcal{M}$ to succeed in the case where $\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} \geq 0$, we must have that $\tau^{l}_{\mathrm{BA}_{\mathcal{M}}} < 0$. As a result

$$\mathbb{P}[(\mathsf{Success}) \cap (\tau^{k}_{\mathrm{AB}_{\mathcal{M}}} \geq 0)] < \mathsf{negl}(n)$$

Thus, from (25)

$$\mathbb{P}[\mathsf{Success}] < \mathsf{negl}(n)$$

Qualitatively, the proof presented here argues that for the adversary to succeed, it needs to either advance the *sync* signal ($\tau_{\mathrm{AB}_{\mathcal{M}}} < 0$), or advance the *response* signal ($\tau_{\mathrm{BA}_{\mathcal{M}}} < 0$). With the use of a secure MAC (or digital signature) and the hardness of superluminal propagation, the adversary can only succeed with a negligible probability.

VI. Secure Constructions

This section specializes the necessary and sufficient conditions for secure clock synchronization to IEEE 1588 PTP. In addition, it presents an alternative to PTP for wireless synchronization—a compliant synchronization system with GNSS-like signals.

A. Secure IEEE 1588 PTP

The necessary and sufficient conditions for secure synchronization, as adapted to IEEE 1588 PTP, are as follows:

1) Stations A and B must use an authenticated encryption scheme to prevent $\mathcal{M}$ from generating valid message/tag pairs.
2) The difference between the path delays between A and B and the shortest possible path delays must be negligible as compared to L. For wireless PTP [31], [32], this implies communicating over the LOS channel as $L \to 0$. For traditional wireline PTP, A and B must attempt to communicate over the (nearly) shortest possible path.
3) The path delay, which is usually estimated from the RTT measurements, must be accurately known *a priori* for secure synchronization. The RTT measurements must be verified against the expected RTT. This implies that the layover time $\bar{\tau}_{\mathrm{BB}}$ must also be known to A.

Note that in the usual PTP formulation, the path delay is measured and used by the time seeker B. To this end, in the usual formulation A sends the transmit time of the *sync* message and the receipt time of the *delay_req* message (in PTP parlance). Similar conventions may be accommodated in the system model presented in this paper, wherein A sends the values of $t^{\mathrm{A}_k}_{\mathrm{A}}$, z^{l}_{A}, and $\bar{\tau}^{kl}_{\mathrm{RTT}}$ to B, and the following calculations may be performed and used at B. However, this would only be a cosmetic change and does not affect the arguments in this paper.

The first security condition has already been proposed in the IEEE 1588-2008 standard. The second condition, however, has not been considered in any of the earlier works in the literature. Following the depiction of *sync* and *response* signal exchange in Fig. 2 and the attack strategy outlined in Section IV-A2, Fig. 3 illustrates an example attack against a PTP implementation that does not satisfy the second necessary condition. Notice that the existence of a shorter time path enables $\mathcal{M}$ to advance the *sync* signal relative to the authentic message from A. Subsequently, $\mathcal{M}$ is able to undo the effect of the advance on the RTT by delaying the *response* signal from B to A. Station A does not measure any abnormality in the RTT, and thus cannot raise an alarm. Meanwhile, synchronization has been compromised at B.

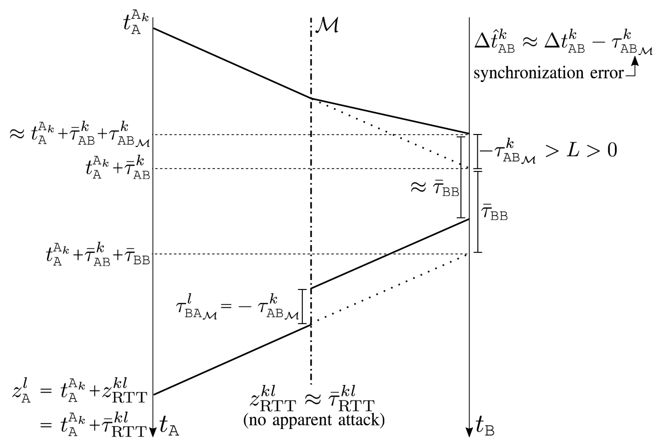

Fig. 3. Illustration of an example attack against a PTP implementation that violates the second necessary condition.

The third condition is similar to the proposal in [11] of measuring the path delays during initialization and monitoring the delays during normal operation. However, [11] requires that B respond to A with zero delay during initialization to enable measurement of the reference delays. This condition is sufficient, but not necessary for secure synchronization. The system is in fact secure even if B is allowed a fixed layover time. Fig. 4 illustrates an example attack against a PTP implementation in violation of the third necessary condition. Note that the uncertainty of the *a priori* estimate of the RTT ($\bar{\sigma}_{\mathrm{RTT}}$) is larger than the alert limit, violating the third necessary condition which requires that the expected RTT be known to within the alert limit (and with much higher accuracy for provable sufficiency). Even though the measured RTT in this case is inconsistent with the expected RTT, it cannot be definitively flagged as an attack since benign variations in the RTT may also have led to the observed RTT.

Interestingly, at first sight, the third security condition in this paper does not resemble the proposed defense in [18] that enforces an upper bound on the synchronization error accumulated between *sync* messages and recommends that B send its timestamps to A periodically for verification. As explained next, this condition is in fact equivalent to the condition of known and measurable RTT, when adapted according to the system model considered in [18].

Note that the requirement of a zero delay in [11], or a short layover time in this paper, enables A to measure the RTT since the transmit time of the lth feature in $s_{\mathtt{B}}$, that is $t_{\mathtt{B}}^{\mathtt{B}_l}$, can be approximately traced back to A's clock to within the alert limit as $t_{\mathtt{A}}^{\mathtt{A}_k} + \bar{\tau}_{\mathtt{AB}}^k + \bar{\tau}_{\mathtt{BB}}$. Enforcing the synchronization error to within

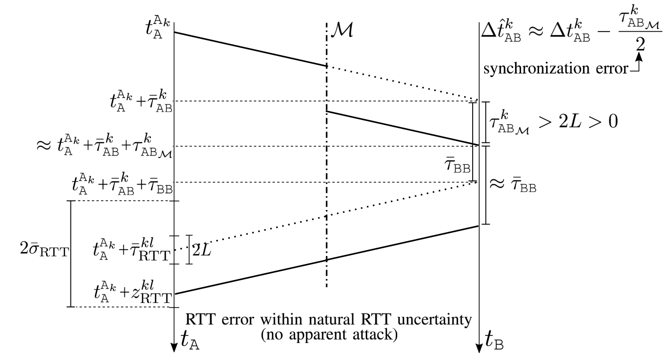

Fig. 4. Illustration of an example attack against a PTP implementation that violates the third necessary condition.

L and transmitting B's timestamp to A achieves the same objective for the defense in [18], since the transmit time from B can be traced back to A's clock with the assumed approximate synchronization. Therefore, the proposed countermeasures in [11] and [18] are two different incarnations of the third security condition proposed in this paper. Of course, the failure of both [11] and [18] to address the second necessary condition makes their proposed defenses vulnerable to an adversary that can communicate along a shorter time path between A and B.

B. Alternative Compliant System

This section describes an alternative wireless clock synchronization protocol that satisfies the set of necessary and sufficient conditions presented in Section IV. The proposed protocol involves bi-directional exchange of GNSS-like pseudo-random codes for continuous clock synchronization, in contrast to discrete packet-based synchronization techniques such as NTP and PTP. It is offered here to illustrate the general applicability of the proposed necessary and sufficient conditions to a range of underlying protocols. Such a protocol can potentially be applied in two-way satellite time transfer and terrestrial wireless clock synchronization systems for continuous clock synchronization, in contrast to the packet-based discrete synchronization in NTP/PTP.

The time master A and the time seeker B communicate wirelessly over the LOS channel between the nodes. To simplify the analysis, it is assumed that A and B securely share long sequences of pseudo-random bits prior to synchronization. These sequences of bits will later enable generation of unpredictable signals. The pseudo-random sequence for A has the form

$$\boldsymbol{b}_{\mathtt{A}} = \left\{ b_{\mathtt{A}}^k \right\}_{k=0}^{N}, \quad b_{\mathtt{A}}^k \in \{0, 1\}$$

The pseudo-random code $C_{\mathtt{A}}(t_{\mathtt{A}})$ for A is then generated as

$$C_{\mathtt{A}}(t_{\mathtt{A}}) = 2b_{\mathtt{A}}^k - 1 \text{ for } t_{\mathtt{A}} \in [t_{\mathtt{A}}^{\mathtt{A}_k}, t_{\mathtt{A}}^{\mathtt{A}_{k+1}}), k \in \{0, 1, 2, \dots\}$$

where $t_{\mathtt{A}}^{\mathtt{A}_k}$ denotes the time according to A at which the start of the kth bit in A's signal is transmitted. The pseudo-random nature of $\boldsymbol{b}_{\mathtt{A}}$ ensures that $C_{\mathtt{A}}(t_{\mathtt{A}})$ has good cross-correlation properties, which enables an accurate measurement of the time-of-arrival of A's signal at B, that is, $\sigma_\epsilon \ll L$. Station

$\mathtt{A}$ modulates a carrier with the code $C_{\mathtt{A}}$ and transmits a signal $s_{\mathtt{A}}(t_{\mathtt{A}})$ whose complex baseband representation is given as

$$s_{\mathtt{A}}(t_{\mathtt{A}}) = C_{\mathtt{A}}(t_{\mathtt{A}}) \exp\left(j\theta_{\mathtt{A}}(t_{\mathtt{A}})\right)$$

This signal is received at $\mathtt{B}$ as

$$\begin{aligned} r_{\mathtt{B}}(t_{\mathtt{A}}, \tau_{\mathtt{AB}}) &= s_{\mathtt{A}}(t_{\mathtt{A}} - \tau_{\mathtt{AB}}) + w_{\mathtt{AB}}(t_{\mathtt{A}}) \\ &= C_{\mathtt{A}}(t_{\mathtt{A}} - \tau_{\mathtt{AB}}) \exp\left(j\theta_{\mathtt{A}}(t_{\mathtt{A}} - \tau_{\mathtt{AB}})\right) + w_{\mathtt{AB}}(t_{\mathtt{A}}) \end{aligned}$$

where all symbols have their usual meanings as detailed in Section III. Station $\mathtt{B}$ captures a window $R^k_{\mathtt{B}}$ of $r_{\mathtt{B}}$ and correlates it with a local replica of $C_{\mathtt{A}}$. The result of the correlation enables $\mathtt{B}$ to detect the start of the kth bit of $C_{\mathtt{A}}$ in the window, and provides a measurement

$$z^k_{\mathtt{B}} = t^{\mathtt{A}_k}_{\mathtt{B}} + w^k_{\mathtt{AB}}$$

of the time-of-arrival of the kth bit at $\mathtt{B}$. Furthermore, the relationship between the start of the kth bit and $t^{\mathtt{A}_k}_{\mathtt{A}}$ enables $\mathtt{B}$ to infer the latter.

If a prior estimate $\bar{\tau}^k_{\mathtt{AB}}$ of $\tau^k_{\mathtt{AB}}$ is available, then $\mathtt{B}$ estimates the clock offset $\Delta t^k_{\mathtt{AB}}$ as in (4).

Similar to the pseudo-random sequence and code construction for $\mathtt{A}$, $\mathtt{B}$ generates its unpredictable code $C_{\mathtt{B}}(t_{\mathtt{B}})$. $\mathtt{A}$ and $\mathtt{B}$ agree on a one-to-one mapping between $C_{\mathtt{A}}$ and $C_{\mathtt{B}}$ such that $\mathtt{B}$ responds with the lth bit of $C_{\mathtt{B}}$ on reception of the start of the kth bit of $C_{\mathtt{A}}$. Furthermore, $\mathtt{A}$ and $\mathtt{B}$ agree that the start of the lth bit of $C_{\mathtt{B}}$ will have a code-phase offset of $\bar{\tau}_{\mathtt{BB}}$ with respect to the start of the kth bit of $C_{\mathtt{A}}$. Station $\mathtt{B}$ transmits the *response* signal as

$$s_{\mathtt{B}}(t_{\mathtt{B}}) = C_{\mathtt{B}}(t_{\mathtt{B}}) \exp\left(j\theta_{\mathtt{B}}(t_{\mathtt{B}})\right)$$

such that

$$t^{\mathtt{B}_l}_{\mathtt{B}} = z^k_{\mathtt{B}} + \bar{\tau}_{\mathtt{BB}}$$

according to the time at $\mathtt{B}$. In true time, the epoch $t^{\mathtt{B}_l}_{\mathtt{B}}$ corresponds to

$$t^{\mathtt{B}_l}_{\mathtt{B}} \rightleftharpoons t^{\mathtt{A}_k}_{\mathtt{A}} + \tau^k_{\mathtt{AB}} + w^k_{\mathtt{AB}} + \tau_{\mathtt{BB}}$$

Station $\mathtt{A}$ receives the *response* as

$$r_{\mathtt{A}} = s_{\mathtt{B}}(t_{\mathtt{B}} - \tau_{\mathtt{BA}}) + w_{\mathtt{BA}}(t_{\mathtt{A}})$$

and captures a window of the signal $R^l_{\mathtt{A}}$. $\mathtt{A}$ correlates $R^l_{\mathtt{A}}$ with a local replica of $C_{\mathtt{B}}$ to detect the start of the lth bit of $C_{\mathtt{B}}$. This enables $\mathtt{A}$ to measure the time-of-arrival

$$z^l_{\mathtt{A}} = t^{\mathtt{B}_l}_{\mathtt{A}} + w^l_{\mathtt{BA}}$$

Moreover, the detection of the lth bit indicates that it was transmitted in response to the receipt of the start of the kth bit of $C_{\mathtt{A}}$. Since $\mathtt{A}$ knows the start time of the kth bit as $t^{\mathtt{A}_k}_{\mathtt{A}}$, it measures the RTT as described in (6).

Note that the exchange of one-time pad sequences enables the proposed system to satisfy the first security condition. Wireless LOS communication satisfies the second security condition, and the knowledge of the code-phase layover offset enables $\mathtt{A}$ to make an accurate prior estimate of the RTT within the alert limit, thereby satisfying the third security condition. Thus, the proposed system complies with all three necessary and sufficient conditions for secure clock synchronization.

VII. System Simulation

This section presents a simulation study of a secure clock synchronization model operating over a simplistic channel model. Unlike the abstract treatment of delays in the security derivations presented earlier, the simulation is carried out with models of delays experienced by the synchronization messages over a real channel. This study also expounds the interplay between slave clock stability, security requirements, attack models, and attack detection thresholds that must be determined in a practical synchronization system. The channel and attack models developed in this simulation are not comprehensive. Rather, relatively simple models are considered to clearly demonstrate the underlying principles. More sophisticated channel and attack models can similarly be analyzed by following the outline of this simulation.

A. Channel Model

The simulated system resembles a traditional local area network, and is schematically depicted in Fig. 5. As before, $\mathtt{A}$ and $\mathtt{B}$ are the time master and seeker stations, respectively. The messages between these stations pass through a series of N routers. Each router is under network traffic loading generated by the nodes labeled $\mathtt{T}$. The routers perform simple packet forwarding, i.e., no cryptographic operations or complex payload modifications are performed. Each router transmits the queued packets at a service rate of 1 Gbps. Each network packet is assumed to have a size of 1542 bytes. The MITM adversary $\mathcal{M}$ maliciously inserts itself along the communication path between $\mathtt{A}$ and $\mathtt{B}$.

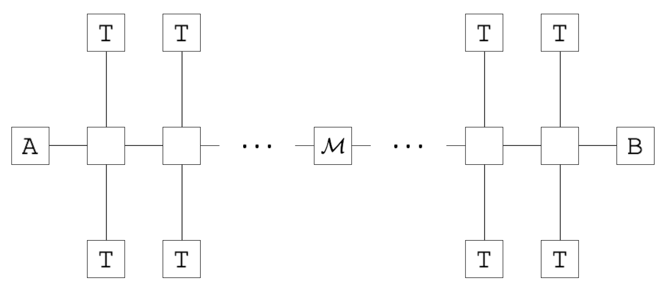

Fig. 5. Schematic diagram of the network topology considered in this section.

The *sync* and *response* packets from $\mathtt{A}$ and $\mathtt{B}$ experience processing and queueing delay at each router, and propagation/link delay between routers. Queueing delay is the duration for which the packet is buffered in the router before it can be transmitted. Processing delay is the time taken by the router to process the packet header, for example, to determine the packet's destination. Since the routers in this simulation perform simple packet forwarding, the processing delay is negligible as compared to the queueing delay [33]. The propagation/link delay is also insignificant for local networks because the propagation speed is a comparable fraction of the speed of light. Thus, only the queueing delay significantly contributes to the overall channel delay variations.

Let the network idle probability for a particular router, denoted by ρ, be defined as the probability of the router queue being empty at a randomly chosen time instant. Since the synchronization packets are delay-sensitive, the routers in this simulation implement non-preemptive priority scheduling for synchronization packets when the queue is not empty. This means that on arrival of a *sync* or *response* packet, the router is allowed to complete the transmission of the data packet currently being serviced, if any, but is required to service the delay-sensitive packet before the other network data in the queue. Since the time period between consecutive *sync-response* pairs is quite large as compared to the RTT for a given pair, it is assumed that a router never has more than one delay-sensitive packet in its queue. Under such scheduling, the delay experienced by the timing messages is best modeled as follows: with probability ρ, the total router delay is zero, and with probability $(1-\rho)$ the total router delay is uniformly distributed between zero and the maximum time to service a packet of length 1542 bytes ($1542 \times 8 \times 2^{-30} \approx 11.49$ microseconds for a Gigabit router).

Given the above channel specifications and values for N and ρ, it is possible to perform a Monte Carlo simulation to obtain the anticipated RTT $\bar{\tau}_{\mathrm{RTT}}$, which is taken to be the empirical mean of the RTT measurements in the simulation, and the associated standard deviation $\bar{\sigma}_{\mathrm{RTT}}$. As shown in Fig. 6, in case of a single *sync-response* pair measurement, the RTT has an empirical mean of 80.34 microseconds and an empirical standard deviation of 17.09 microseconds with $N = 10$ and $\rho = 0.3$. Observe that even for a relatively small N, the empirical distribution approaches the Gaussian shape, but has slightly heavier tail on the higher end of the delay. The distribution for mean of batches of 10 observations has a smaller empirical standard deviation of 5.41 microseconds.

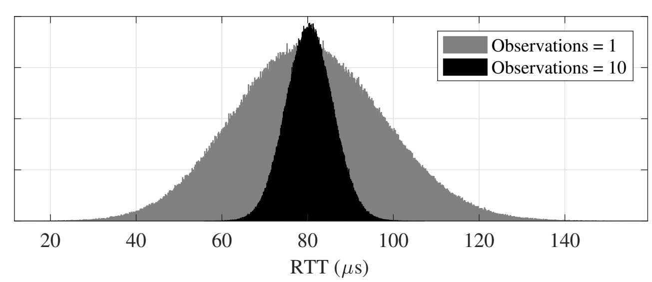

Fig. 6. Empirical distribution of the RTT of *sync-response* pairs through a network of $N = 10$ routers with network idle probability of $\rho = 0.3$. The light-shaded histogram shows the empirical distribution of the RTT of a single *sync-response* pair. The dark-shaded histogram shows the corresponding distribution for the mean of batches of 10 observations of the RTT.

B. System and Security Requirements

The clock at the time seeker $\mathtt{B}$ drifts with respect to the true time clock at $\mathtt{A}$ unless corrected by a *sync* message from $\mathtt{A}$. As before, let L denote the alert limit for the system. Let T denote a time duration over which a perfectly synchronized clock at $\mathtt{B}$ at the beginning of the duration, absent an adversary, drifts more than $L_{\mathcal{N}}$ for some $L_{\mathcal{N}} < L$ with a probability smaller than an acceptably small bound $\mathbb{P}_{\epsilon}$.

In the system under simulation, the clock offset for $\mathtt{B}$ is estimated and corrected for every T seconds. By definition of T, it holds that if the clocks at $\mathtt{A}$ and $\mathtt{B}$ are perfectly synchronized after every T seconds, then the natural drift envelope of $\mathtt{B}$'s clock does not exceed L with an unacceptably high probability. Define

$$L_{\mathcal{M}} \triangleq L - L_{\mathcal{N}}$$

Observe that if an adversary is able to introduce a synchronization error larger than $L_{\mathcal{M}}$, then the system is compromised since the natural drift of the clock at $\mathtt{B}$ could potentially lead to a clock offset greater than L before the next synchronization interval, with a probability greater than $\mathbb{P}_{\epsilon}$. Thus, $\mathtt{A}$ must flag any adversarial delay greater than $L - L_{\mathcal{N}}$ with probability higher than a desired detection probability, denoted by $\mathbb{P}_{\mathrm{D}}$. It is worth noting that this practical complication of the magnitude of $L_{\mathcal{N}}$ was abstracted in the sufficiency proof, where the threshold was set to $L - \delta$ for $\delta > 0$.

In general, $\mathtt{A}$ makes multiple measurements of the RTT between $\mathtt{A}$ and $\mathtt{B}$ over the time period T. As shown in Fig. 6, the mean of multiple observations over T has a distribution with a smaller standard deviation as compared to that of a single observation. In the simulated system, if no attack is detected, $\mathtt{A}$ updates $\bar{\tau}^{k}_{\mathrm{AB}}$ every T seconds based on the empirical mean of the RTT measurements made over that period. Note that even though $\bar{\tau}^{k}_{\mathrm{AB}}$ is updated based on the measurements, no updates are applied to $\bar{\tau}_{\mathrm{RTT}}$ and $\bar{\sigma}_{\mathrm{RTT}}$, which are pre-determined by simulation or measurements under a secure calibration campaign.

The empirical mean of the measured RTT is taken as the test statistic to detect an attack. For the attack model detailed next, it can be shown that this test statistic becomes optimal for large values of N [34].

C. Attack Model

The synchronization system considered in this simulation complies with the necessary security conditions presented in this paper. Consequently, the adversary $\mathcal{M}$ is unable to advance the *sync* or *response* messages, and can only increase the RTT measured by $\mathtt{A}$ relative to the authentic RTT. This simulation considers a simple adversary model that introduces a fixed delay in the measured RTT. In order to conceal its presence while compromising synchronization with appreciable probability, $\mathcal{M}$ introduces a delay of $L_{\mathcal{M}} + \xi$ seconds for some small $\xi > 0$.

Let H_0 denote the null hypothesis (no attack), and H_1 denote the alternative hypothesis. Under H_0, the measured RTT at $\mathtt{A}$ is drawn from the distribution that was used to calibrate/simulate the channel delay distribution, while under H_1, the measured RTT is drawn from a distribution that is shifted from the calibration distribution by $L_{\mathcal{M}} + \xi$. This is visually depicted in Fig. 7. Given a detection threshold λ, the dark-shaded region in Fig. 7 denotes the probability of false alarm, $\mathbb{P}_{\mathrm{F}}$, while the light-shaded region denotes the probability of missed detection $(1-\mathbb{P}_{\mathrm{D}})$. In observing Fig. 7, it might be argued, and holds true, that a reasonable attacker may introduce noise in the introduced delay to inflate the width of the distribution under H_1 and thereby decrease the probability of detection of an attack. However, in that case, the empirical mean test statistic is no longer optimal. Instead, $\mathtt{A}$ would incorporate the observed variance of the RTT in its test statistic in addition to the emprical mean. In short, the attack model in this simulation is not comprehensive, as explained previously. For a more sophisticated treatment of sensor deception and protection techniques, the reader may refer to [35].

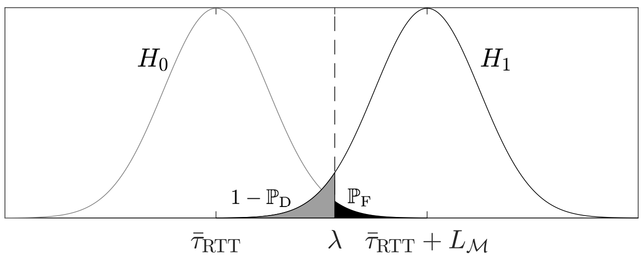

Fig. 7. Representation of the distributions under H_0 and H_1 along with the detection threshold and the associated $\mathbb{P}_\mathrm{F}$ and $\mathbb{P}_\mathrm{D}$.

D. Simulation

The system and attack described above have been simulated with $N = 10$ and $\rho = 0.3$ for all routers. The adversarial delay $L_\mathcal{M} + \xi$ is set to 10 microseconds, and the required probability of detection $\mathbb{P}_\mathrm{D}$ is set to 0.999. The number of RTT observations made in time T are varied between 1 and 200. Given the number of observations, and a required $\mathbb{P}_\mathrm{D}$, the system is simulated under H_1 for 10^6 detection epochs and the maximum possible detection threshold λ that satisfies the detection probability is obtained. Subsequently, the system is simulated under H_0 and the number of test statistics exceeding the threshold λ are recorded. The frequency of such epochs is reported as the probability of false alarm $\mathbb{P}_\mathrm{F}$.

Fig. 8 shows the above procedure for 80 RTT measurements made per test statistic. In this case, λ is obtained to be 84.53 microseconds and the corresponding $\mathbb{P}_\mathrm{F}$ is 1.59%. Fig. 9 shows a log-log plot of $\mathbb{P}_\mathrm{F}$ as a function of the number of observations made per test statistic. When the number of observations is greater than 160, no false alarms were observed with 10^6 trials. For the given channel delay variation statistics, the probability of false alarm is very high for small number of observations per decision epoch since the threshold λ that must be set to detect an attack with the required $\mathbb{P}_\mathrm{D}$ is large in comparison to the minimal delay that the adversary must introduce to compromise synchronization ($L_\mathcal{M}$). For a more stable channel, such as a wireless or PTP-aware channel, fewer measurements per decision epoch would suffice.

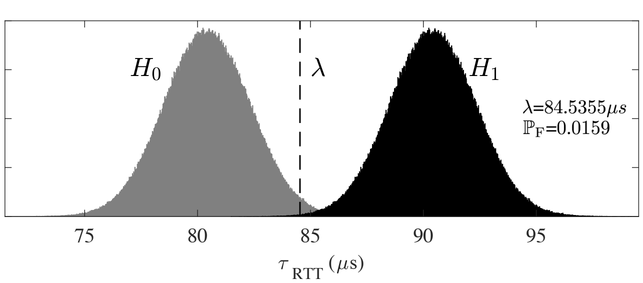

Fig. 8. Distribution of the test statistic under H_0 and H_1 for 80 RTT measurements per decision epoch. ($N = 10$, $\rho = 0.3$, $L_\mathcal{M} + \xi = 10\mu$s, $\mathbb{P}_\mathrm{D} = 0.999$)

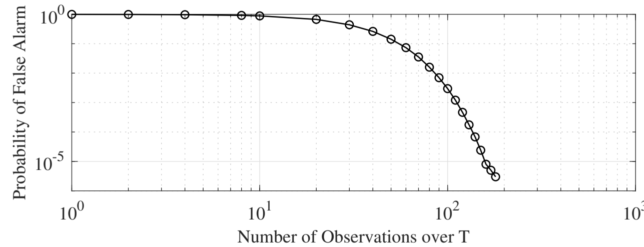

Fig. 9. Probability of false alarm as a function of number of observations per decision epoch. ($N = 10$, $\rho = 0.3$, $L_\mathcal{M} + \xi = 10\mu$s, $\mathbb{P}_\mathrm{D} = 0.999$)

E. Practical Implications

Section V-A makes fairly remarkable assumptions about the synchronization system to show provably secure time transfer. For instance, it requires that errors in the *a priori* estimate of the RTT of the timing messages be negligible compared to the alert limit. Nonetheless, as shown in this section, for a given channel with bounded delay variations and a given slave clock, some level of security guarantee can be made for a synchronization system that satisfies the necessary and sufficient conditions presented herein. For concreteness, consider a system that requires an alert limit of $L = 100$ microseconds and a slave clock that drifts no more than $L_\mathcal{N} = 50$ microseconds over a period of $T = 1$ second with acceptably high probability $(1 - \mathbb{P}_\epsilon)$. Then, for $N = 10$ and $\rho = 0.3$, if $\mathtt{A}$ makes 10 RTT measurements over 1 second, the empirical mean test statistic is distributed as the dark-shaded distribution in Fig. 6 with a standard deviation of ≈ 5.4 microseconds. For $L_\mathcal{M} = L - L_\mathcal{N} = 50$ microseconds, a threshold of $\approx \bar{\tau}_\mathrm{RTT} + 30$ microseconds will yield a missed detection rate of approximately 1 in 15000, and a false alarm rate of approximately 3.5 in 1 million. With a more stable slave clock or more measurements per second, these probabilities can be made more favorable.

Another important concern that has not been addressed in the simulation is that of the incorporation of cryptographic constructs in the synchronization protocol. The encryption and decryption algorithms are often complex and take non-negligible processing time to execute. However, note that at $\mathtt{A}$, the *sync* message is timestamped *after* the encryption process, and thus the time taken for encryption is inconsequential. At $\mathtt{B}$, it is important to concede that the decryption of the *sync* message and the encryption of the *response* message cannot be assumed to happen instantaneously. This has been accounted for by allowing the layover time $\bar{\tau}_\mathrm{BB}$ for the cryptographic processes to execute. Once again, the receipt timestamp of the *response* message at $\mathtt{A}$ is applied *before* the decryption process, and hence the decryption time at $\mathtt{A}$ is inconsequential. Thus, compliance with the first security condition must not pose significant practical challenges.

VIII. Conclusions

A fundamental theory of secure clock synchronization was developed for a generic system model. The problem of secure clock synchronization was formalized with explicit assumptions, models, and definitions. It was shown that all possible one-way clock synchronization protocols are vulnerable to replay attacks. A set of necessary conditions for secure two-way clock synchronization was proposed and proved. Compliance with these necessary conditions with strict upper bounds was shown to be sufficient for secure clock synchronization, which is a significant result for provable security in critical infrastructure. The general applicability of the set of security conditions was demonstrated by specializing these conditions to design a secure PTP protocol and an alternative secure two-way clock synchronization protocol with GNSS-like signals. Results from a simulation with models of channel delays were presented to expound the interplay between slave clock stability, security requirements, attack models, and attack detection thresholds.

ACKNOWLEDGMENTS

This project has been supported by the National Science Foundation under Grant No. 1454474 (CAREER), by the Data-supported Transportation Operations and Planning Center (DSTOP), a Tier 1 USDOT University Transportation Center, and by the U.S. Department of Energy under the TASQC program led by Oak Ridge National Laboratory.

REFERENCES

[1] A. Phadke, B. Pickett, M. Adamiak, M. Begovic, G. Benmouyal, R. Burnett Jr, T. Cease, J. Goossens, D. Hansen, M. Kezunovic *et al.*, "Synchronized sampling and phasor measurements for relaying and control," *IEEE Transactions on Power Delivery*, vol. 9, no. 1, pp. 442–452, 1994.

[2] J. G. McNeff, "The global positioning system," *IEEE Transactions on Microwave Theory and Techniques*, vol. 50, no. 3, pp. 645–652, 2002.

[3] J. J. Angel, "When finance meets physics: The impact of the speed of light on financial markets and their regulation," *Financial Review*, vol. 2, no. 49, pp. 271–281, 2014.

[4] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. J. Furman, S. Ghemawat, A. Gubarev, C. Heiser, P. Hochschild *et al.*, "Spanner: Googles globally distributed database," *ACM Transactions on Computer Systems (TOCS)*, vol. 31, no. 3, p. 8, 2013.

[5] L. D. Shapiro, "Time synchronization from Loran-C," *IEEE Spectrum*, vol. 8, no. 5, pp. 46–55, 1968.

[6] A. Bauch, P. Hetzel, and D. Piester, "Time and frequency dissemination with DCF77: From 1959 to 2009 and beyond," *PTB-Mitteilungen*, vol. 119, no. 3, pp. 3–26, 2009.

[7] D. W. Allan and M. A. Weiss, *Accurate time and frequency transfer during common-view of a GPS satellite*. Electronic Industries Association, 1980.

[8] M. L. Psiaki and T. E. Humphreys, "GNSS spoofing and detection," *Proceedings of the IEEE*, vol. 104, no. 6, pp. 1258–1270, 2016.

[9] K. D. Wesson, J. N. Gross, T. E. Humphreys, and B. L. Evans, "GNSS signal authentication via power and distortion monitoring," *IEEE Transactions on Aerospace and Electronic Systems*, 2018, to be published; preprint available at https://arxiv.org/abs/1702.06554.

[10] D. Chou, L. Heng, and G. Gao, "Robust GPS-based timing for phasor measurement units: A position-information-aided approach," in *Proceedings of the ION GNSS+ Meeting*, 2014.

[11] M. Ullmann and M. Vögeler, "Delay attacks — Implication on NTP and PTP time synchronization," in *Precision Clock Synchronization for Measurement, Control and Communication, 2009. ISPCS 2009. International Symposium on*. IEEE, 2009, pp. 1–6.

[12] T. Mizrahi, "A game theoretic analysis of delay attacks against time synchronization protocols," in *Precision Clock Synchronization for Measurement Control and Communication (ISPCS), 2012 International IEEE Symposium on*. IEEE, 2012, pp. 1–6.

[13] B. Moussa, M. Debbabi, and C. Assi, "A detection and mitigation model for PTP delay attack in an IEC 61850 substation," *IEEE Transactions on Smart Grid*, 2016.

[14] Q. Yang, D. An, and W. Yu, "On time desynchronization attack against IEEE 1588 protocol in power grid systems," in *Energytech, 2013 IEEE*. IEEE, 2013, pp. 1–5.

[15] J.-C. Tournier and O. Goerlitz, "Strategies to secure the IEEE 1588 protocol in digital substation automation," in *Critical Infrastructures, 2009. CRIS 2009. Fourth International Conference on*. IEEE, 2009, pp. 1–8.

[16] S. Bhamidipati, Y. Ng, and G. X. Gao, "Multi-receiver GPS-based direct time estimation for PMUs," in *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016), Portland, OR*, 2016.

[17] Y. Ng and G. X. Gao, "Robust GPS-based direct time estimation for PMUs," in *Position, Location and Navigation Symposium (PLANS), 2016 IEEE/ION*. IEEE, 2016, pp. 472–476.

[18] R. Annessi, J. Fabini, and T. Zseby, "SecureTime: Secure multicast time synchronization," *arXiv preprint arXiv:1705.10669*, 2017.

[19] L. Narula and T. E. Humphreys, "Requirements for secure wireless time transfer," in *Proceedings of the IEEE/ION PLANS Meeting*, Savannah, GA, 2016.

[20] R. C. Merkle, "Secure communications over insecure channels," *Communications of the ACM*, vol. 21, no. 4, pp. 294–299, 1978.

[21] C. H. Bennett and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing," 1984.

[22] A. K. Ekert, "Quantum cryptography based on Bell's theorem," *Physical review letters*, vol. 67, no. 6, p. 661, 1991.

[23] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance*, revised second ed. Lincoln, Massachusetts: Ganga-Jumana Press, 2012.

[24] K. D. Wesson, M. P. Rothlisberger, and T. E. Humphreys, "A proposed navigation message authentication implementation for civil GPS anti-spoofing," in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2011.

[25] D. M. Akos, "Who's afraid of the spoofer? GPS/GNSS spoofing detection via automatic gain control (AGC)," *Navigation, Journal of the Institute of Navigation*, vol. 59, no. 4, pp. 281–290, 2012.

[26] T. E. Humphreys, *Springer Handbook of Global Navigation Satellite Systems*. Springer, 2017, ch. Interference, pp. 469–504.

[27] M. Meurer, A. Konovaltsev, M. Cuntz, and C. Hättich, "Robust joint multi-antenna spoofing detection and attitude estimation using direction assisted multiple hypotheses RAIM," in *Proceedings of the 25th Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2012)*. ION, 2012.

[28] D. Borio, "PANOVA tests and their application to GNSS spoofing detection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 1, pp. 381–394, Jan. 2013.

[29] M. Bellare and C. Namprempre, "Authenticated encryption: Relations among notions and analysis of the generic composition paradigm," *Advances in Cryptology–ASIACRYPT 2000*, pp. 531–545, 2000.

[30] S. Goldwasser, S. Micali, and R. L. Rivest, "A digital signature scheme secure against adaptive chosen-message attacks," *SIAM Journal on Computing*, vol. 17, no. 2, pp. 281–308, 1988.

[31] A. Mahmood, G. Gaderer, H. Trsek, S. Schwalowsky, and N. Kerö, "Towards high accuracy in IEEE 802.11 based clock synchronization using PTP," in *Precision Clock Synchronization for Measurement Control and Communication (ISPCS), 2011 International IEEE Symposium on*. IEEE, 2011, pp. 13–18.

[32] T. Cooklev, J. C. Eidson, and A. Pakdaman, "An implementation of IEEE 1588 over IEEE 802.11b for synchronization of wireless local area network nodes," *IEEE Transactions on Instrumentation and Measurement*, vol. 56, no. 5, pp. 1632–1639, Oct. 2007.

[33] R. Ramaswamy, N. Weng, and T. Wolf, "Characterizing network processing delay," in *Global Telecommunications Conference, 2004. GLOBECOM'04. IEEE*, vol. 3. IEEE, 2004, pp. 1629–1634.

[34] H. L. Van Trees, *Detection, estimation, and modulation theory, part I: detection, estimation, and linear modulation theory*. John Wiley & Sons, 2004.

[35] J. Bhatti and T. Humphreys, "Hostile control of ships via false GPS signals: Demonstration and detection," *Navigation, Journal of the Institute of Navigation*, vol. 64, no. 1, 2017.