8 Synchronization

This chapter deals with another major challenge connected to CoMP, namely the synchronization of cooperating and cooperatively served devices in time and frequency. On one hand, there are different local oscillators in each base station and mobile terminal that lead to deviations in the carrier frequency according to its nominal value. On the other hand, there are variations in the symbol timing between each transmitter and receiver station. Both effects need to be compensated by synchronization techniques.

In cellular networks, we can distinguish between a network synchronization among all involved base stations and the alignment of the user equipments to that time and frequency reference. The basic definitions of the synchronization terms as well as procedures for the reference network synchronization are described in Section 8.1. The impact of symbol timing mismatches on CoMP is then treated in Section 8.2, before Section 8.3 concludes this chapter with the analysis of the impact of residual carrier frequency offsets on CoMP performance.

8.1 Synchronization Concepts

D. Richard Brown III and Andrew G. Klein

Synchronization is the process of establishing a common notion of time among two or more entities. In the context of wired and wireless communication networks, synchronization enables coordination among the nodes in the network and can facilitate applications such as distributed sensing. Precise synchronization can also facilitate scheduling of communication resources as well as interference avoidance in multi-access networks. This section provides an overview of some of the synchronization concepts and techniques used in coordinated communication networks.

8.1.1 Synchronization Terminology

In the context of wireless communication networks, each node in the network keeps a local notion of time, i.e. a clock, by counting cycles of a local oscillator (LO). Among other parameters, all oscillators are characterized in terms of their nominal frequency and accuracy. The accuracy of an oscillator is typically specified in parts per million (PPM) with respect to the nominal frequency. For example, a 10 MHz LO with ±10 PPM accuracy oscillates with a frequency between ±100 Hz of the LO's nominal 10 MHz frequency. Low-cost oscillators typically provide ±100 PPM accuracy [FMd09], and higher-cost temperature-compensated or oven-controlled oscillators can provide accuracies better than ±1 PPM [Rak0[illegible] clock derived from an unsynchronized low-cost ±100 PPM LO can gain or loose, with respect to a perfect reference clock, up to 8.64 seconds in a day.

Two clocks are said to be perfectly *syntonized* if they agree exactly on the duration of an interval between two events. In other words, syntonized clocks share the same rate or frequency, but there is no requirement for the clocks to agree on the time of a single event. Syntonized clocks are sometimes said to be frequency synchronized. A clock can be said to be syntonized to a specified level of uncertainty if the frequency difference with respect to a reference clock (often normalized and specified in PPM) is no more than the uncertainty. This frequency difference is commonly called *frequency offset* or *skew* and is typically specified in statistical terms, e.g. standard deviation or maximum frequency offset. For example, the LTE specification requires user equipments (UEs) to have a maximum frequency offset of $\pm 0.1 ppm$ [STE09].

Two clocks are said to be perfectly *synchronized* if they agree exactly on the time of occurrence of an event at an arbitrary time. Note that synchronized clocks must also be syntonized since unsyntonized clocks can only agree on the occurrence of an event at a particular time. Clocks can also be said to be synchronized to a specified level of uncertainty if they do not agree precisely on the time of occurrence of an event, but the difference in the measured event times between the clocks is no more than the uncertainty. The time difference between two clocks is commonly called *clock offset* or *phase offset* and is typically specified in statistical terms, e.g. standard deviation or maximum clock offset.

Since each node in a wireless network keeps time with its own LO, syntonation and synchronization is necessary to establish a common time base among the nodes in the network to a desired level of precision. In fact, if the nodes in the network have a maximum clock offset requirement, periodic frequency and phase re-synchronization is necessary to correct for unavoidable phase drift between any pair of nodes caused by frequency offset and oscillator instabilities. Several factors can affect the accuracy of phase and frequency synchronization among nodes in communication networks. These factors include local oscillator stability, network stability, and the re-synchronization interval, i.e. how often synchronization messages are exchanged among the nodes in the network.

Figure 8.1 shows an example of clock and frequency offset between two clocks labeled "clock A" and "clock B". The re-synchronization interval in this example is denoted as T. Prior to the first synchronization attempt at time t_1, the phase of clock B is ahead of clock A and the phases of the clocks are drifting apart

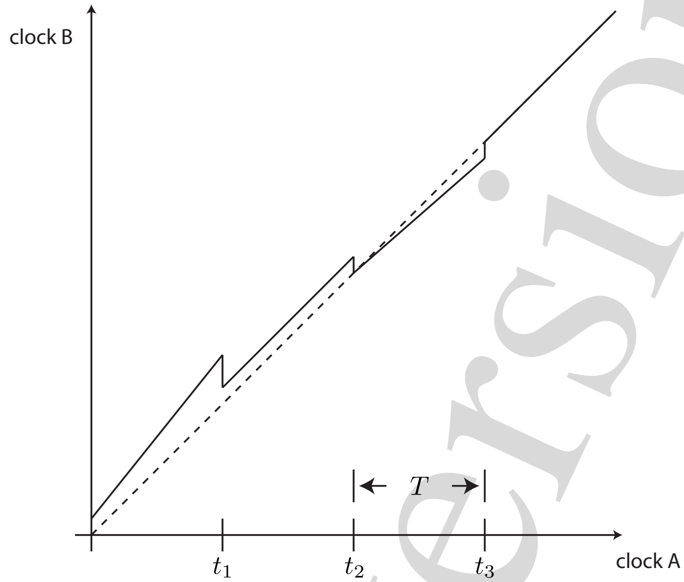

Figure 8.1 Clock and frequency offset example.

due to the frequency offset. At time t_1, clock B adjusts its frequency and phase in response to a synchronization attempt with clock A. In this example, the clocks become syntonized at t_1, but they remain unsynchronized. At time t_2, the clocks again attempt to synchronize. In this example, clock B and clock A are temporarily synchronized at time t_2, but become unsyntonized. The nodes become synchronized and syntonized after the synchronization attempt at t_3.

In the following sections, we provide an overview of several synchronization techniques suitable for distributing a common notion of time to nodes in wired and wireless communication networks. We begin with a discussion of two *network synchronization* techniques: network time protocol (NTP) and precision time protocol (PTP), i.e. IEEE 1588. We then describe *satellite-based synchronization* techniques with a focus on global positioning system (GPS). We conclude with a discussion of *endogenous wireless distributed synchronization* techniques suitable for the precisely synchronizing and syntonizing the carriers of wireless transmitters for distributed phase coherent communication.

8.1.2 Network Synchronization

The Network Time Protocol

The NTP is a protocol for synchronizing the clocks of nodes that are connected through variable-latency networks [Mil91]. NTP is an application-layer protocol that operates over the Internet protocol (IP), and can therefore be implemented completely in software. The protocol has been in use since the 1980's, and today it is responsible for synchronizing the clocks of the majority of computers con-

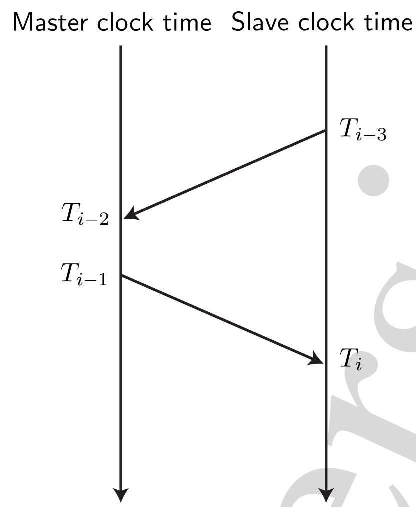

Figure 8.2 NTP message exchange.

nected to the Internet. Nodes in the network are assigned to a class or *stratum*, and those with the lowest stratum number are assumed to be perfectly synchronized with Coordinated Universal Time (UTC). Nodes with higher stratum numbers synchronize their clocks with nodes having lower stratum numbers. This hierarchical structure of NTP results in it being highly scalable.

To estimate clock offsets, a master and slave exchange timestamps which are 64-bit descriptions their current local clock time. Figure 8.2 demonstrates the exchange of timestamps between a master and slave. If T_i, T_{i-1}, T_i, and T_{i-1} are the four most recent timestamps, then the clock offset of the slave relative to the master at time T_i can be calculated via

$$\Delta_i = \frac{T_{i-2} + T_{i-1} - T_{i-3} - T_i}{2}. \tag{8.1}$$

Since each NTP message contains the last three timestamps $T_{i-1}, T_{i-2}, T_{i-3}$, and the final timestamp T_i is estimated upon arrival of the message, the clock offset can be estimated from a single message exchange between slave and master.

Equation (8.1) implicitly assumes that the two transmission paths are symmetric and have equal delay. In practice, however, network delays are stochastic quantities. Consequently, NTP performs multiple offset estimates in combination with a filtering and selection scheme to obtain a more accurate estimate of the clock offset. The estimated clock offsets are fed to a Type-II adaptive parameter phase-locked loop (PLL), which corrects the LO phase and frequency. An adaptive Type-II PLL has one integrator in the loop filter (or two poles in the open-loop transfer function) and continuously adjusts the phase and frequency [Smi86].

The accuracy of the protocol depends on a variety of factors, including the update interval and network topology. Several studies (e.g. [Mil03, MTH97, KZM07, Min99]) have investigated the performance of NTP under typical use, showing that clock offsets have a standard deviation on the order of several milliseconds, and residual frequency offsets on the order of ± 0.1 PPM.

The Precision Time Protocol

Also known as IEEE 1588 [IEE08a], the PTP attains sub-microsecond accuracy which is necessary in applications such as networked control systems and precision machinery in factories. The phase and frequency correction in PTP are quite similar in principle to NTP: after a sequence of messages are exchanged between slave and master, the clock offsets are estimated through filtering and selection, and are used to adjust a PLL which corrects the LO phase and frequency. There are, however, several fundamental differences between PTP and NTP. The primary difference is that PTP is implemented in hardware rather than software. By moving the clock synchronization as close to the physical layer as possible, sources of jitter and processing delay introduced in network layers higher up the stack can be mitigated. In addition, PTP is primarily intended to be used in a local area network (LAN) setting as opposed to NTP, which may synchronize to an Internet clock reference located some far distance away. While PTP can achieve a greater accuracy than NTP, it does require the use of dedicated hardware. Performance of PTP will again depend on a variety of factors, including the quality of the LO, as well as the network topology. Products already available on the market today [Sem10] claim clock offsets within 1 μs and frequency offsets better than ± 0.01 PPM. Similar results were achieved [Ton05] in a test of PTP over a metropolitan area network.

The most recent version of the standard, referred to as IEEE 1588-2008, offers a *transparent clock* mode which requires dedicated network switches that support the standard. Such switches employ a transparent clock that further minimizes delay by providing an alternate local clock for network nodes so that they need not rely on the master clock. This mode permits maximum clock offset errors on the order of tens of nanoseconds [HJ10].

8.1.3 Satellite-Based Synchronization

A Global Navigation Satellite System (GNSS) permits nodes to determine their location to within a few meters using time signals received line-of-sight from satellites. While the primary intent of a GNSS is for determining position information, such systems are also very useful as an accurate, common clock reference. In contrast to NTP and PTP, clock synchronization using a GNSS is done wirelessly using one-way communication links (i.e. by receiving signals broadcast from the satellites). In order for a terrestrial node to be able to receive the relatively weak signals from distant satellites, however, a line-of-sight link is typically necessary.

In the absence of precise location information, a node must be able to receive signals from four satellites since there are four unknowns: latitude, longitude, altitude, and time. If precise location information is available, only one satellite is needed for clock synchronization since propagation delay is known.

Examples of GNSS's include the United States' GPS, the Russian GLObal NAvigation Satellite System (GLONASS), and the European Galileo system. As stated in [LAK99], GPS provides clock synchronization "to better than 100 ns in time and 10^{-13} in frequency." Other satellite systems are expected to give synchronization accuracy of a similar order, as they share many of the same parameters as GPS [HP05a].

8.1.4 Endogenous Distributed Wireless Carrier Synchronization

Coherent downlink CoMP techniques have recently been proposed in which the base stations (BSs) transmit with phase-aligned carriers such that the *bandpass signals* are aligned at the intended destination. These coherent transmission techniques require the BSs to accurately pre-compensate for the downlink channel phases and maintain close synchronization. One approach to coherent downlink CoMP, as discussed in Section 13.3, is to closely syntonize the BSs' carriers using, for example, highly stable GPS-referenced local oscillators. Coherent downlink transmission can then be achieved by having the mobiles estimate the downlink channel state information (CSI) and feeding back this CSI to the BSs for carrier phase pre-compensation. However, periodic downlink CSI re-estimation and low-latency feedback on the order of a few milliseconds is necessary to maintain phase coherence at the mobile.

A different approach to coherent downlink CoMP is to have the BSs endogenously synchronize their carriers without the aid of GPS and without CSI feedback from the mobiles. The "two-way" downlink beamforming technique proposed in [PD10] is one example of this approach. Two-way downlink beamforming is a retrodirective transmission technique in which the BSs closely synchronize their carriers (in both phase and frequency) to emulate a conventional retrodirective antenna array [Pon64] and achieve coherent downlink transmission through uplink channel conjugation. Note that this approach requires uplink/downlink channel reciprocity and also precise synchronization of the BSs.

To understand just how accurately the BSs must be synchronized to facilitate retrodirective transmission, consider the two-transmitter distributed beamforming scenario shown in Fig. 8.3 where both BSs simultaneously transmit unmodulated carriers at radian frequency ω with the goal of having the carriers arrive with identical phase, i.e. "coherently combine", at the destination, i.e. the mobile. Note that the BSs are implicitly syntonized in this scenario, but they are unsynchronized such that transmitter 2 has a clock offset of Δ with respect to transmitter 1. After propagation through the unit-gain single-path channels,

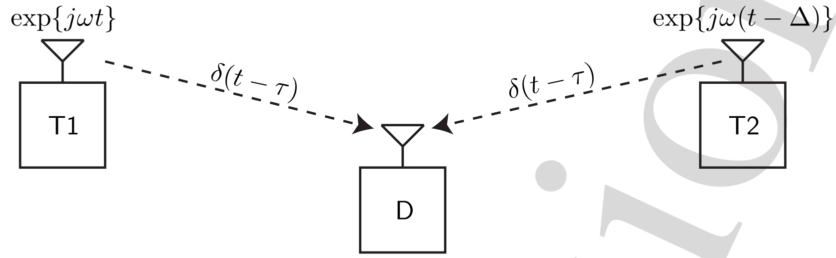

Figure 8.3 Two-transmitter one-destination distributed beamforming scenario.

the received signal at the destination can be written as

$$y(t) = \exp\{j\omega(t-\tau)\} + \exp\{j\omega(t-\Delta-\tau)\}$$

where the baseband signals modulated by each carrier are omitted for clarity. The received power can be computed as $|y(t)|^2 = 2 + 2\cos(\omega\Delta) \leq 4$. When the transmitters are perfectly synchronized, i.e. $\Delta = 0$, the carriers combine coherently at the destination and the received power $|y(t)|^2 = 4$. This corresponds to the “ideal coherent” case in distributed beamforming. When the transmitters are not synchronized, the received power will be less than in the ideal coherent case. To illustrate the effect of unsynchronized transmitters, the clock offset Δ can be modeled as a zero-mean Gaussian distributed random variable with standard deviation σ_Δ. Fig. 8.4 shows the received power at the destination as a function of σ_Δ and carrier frequency $f_0 := \frac{\omega}{2\pi}$. This example shows that, even at the lowest carrier frequency of 800 MHz, the standard deviation of the transmitter clock offset must be smaller than approximately 130 picoseconds in order to achieve, on average, 90% or better of the ideal coherent received power.

Since conventional synchronization techniques like GPS and PTP are unable to provide the accuracy required for retrodirective downlink transmission at typical radio frequencies, several recent studies have focused on the development of precise endogenous distributed wireless carrier synchronization techniques including full-feedback closed-loop [TP02], one-bit closed-loop [MHMB05, MWMR06, MHMBew], master-slave open-loop [MBM07], round-trip open-loop carrier synchronization [DPM05, DH08], and two-way open-loop carrier synchronization [PD10]. Each of these techniques has advantages and disadvantages in particular applications, as discussed in the survey article [MDMH09].

Many of the distributed wireless carrier synchronization techniques described in [MDMH09] operate on the principle of exchanging beacons between the BSs and synchronizing carriers based on estimates of the phase and frequency of these beacons. For example, in the two-way carrier synchronization protocol [PD10], a series of “forward beacons” are exchanged from node 1 to node 2 and so on to node M where node m transmits a periodic extension of the beacon it received from node $m-1$. A series of “backward beacons” are also exchanged in the same way from node M to node $M-1$ and so on to node 1. Each node then sums its

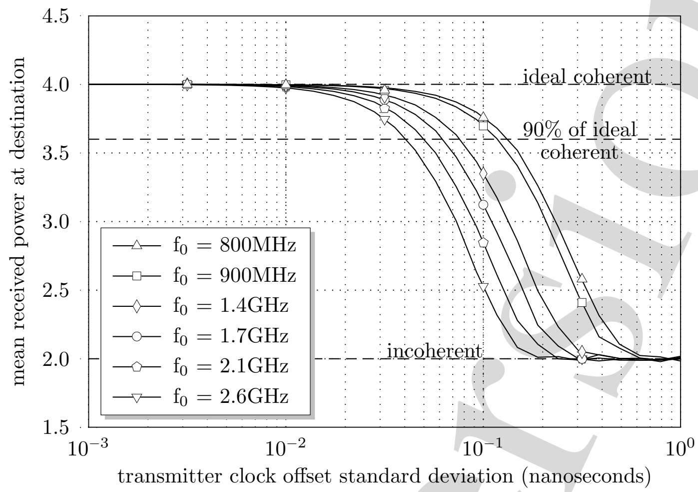

Figure 8.4 The effect of transmitter clock offset on distributed beamforming power for several common cellular carrier frequencies.

phase and frequency estimates obtained from the forward and backward beacons to derive a synchronized local oscillator with frequency and phase identical (in the absence of estimation error) to the other nodes in the system. These synchronized local oscillators can then be used to enable retrodirective downlink transmission from two or more BSs to a mobile in the network.

While the various carrier synchronization techniques described in [MDMH09] differ in terms of how beacons are exchanged, the overall performance of each of these techniques tends to be limited by the accuracy of the beacon phase and frequency estimators. A common technique for the estimation of the phase and frequency of a single tone in additive white Gaussian noise is the maximum likelihood estimator (MLE) [RB74]. Under mild regularity conditions [H.V94], the MLE is known to asymptotically achieve the Cramér-Rao lower bound (CRLB). Given a beacon of amplitude a in complex Gaussian white noise with power spectral density $\frac{N_0}{2}$, the CRLB is given as [RB74]

$$\operatorname{cov}\left\{\left[\tilde{\omega}, \tilde{\theta}\right]^{\top}\right\} \succeq \frac{N_0}{a^2}\begin{bmatrix} \frac{12}{T^3} & \frac{-6}{T^2} \\ \frac{-6}{T^2} & \frac{4}{T} \end{bmatrix}$$

where T is the duration of the observation, and the notation $\mathbf{A} \succeq \mathbf{B}$ means that $\mathbf{A} - \mathbf{B}$ is positive semi-definite.

As an example, Fig. 8.5 shows the clock and frequency offset standard deviations for the two-way carrier synchronization protocol developed in [PD10]. This example assumes seven transmitters serially exchange 1 GHz wireless beacons according to the two-way synchronization protocol. Each beacon has a duration $T = 1$ ms and a total of 12 beacons are exchanged to synchronize the transmitters. The transmitters use MLE to form local phase and frequency estimates

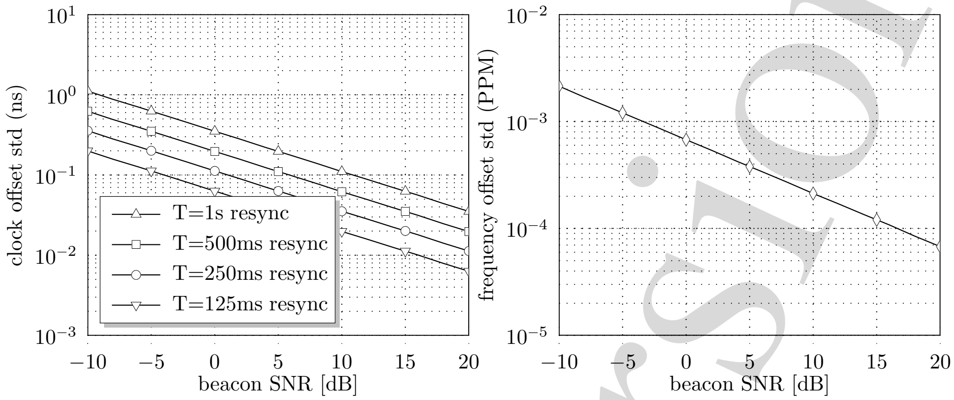

(a) Clock offset std. deviation. (b) Frequency offset std. deviation.

Figure 8.5 Clock and frequency offset standard deviations as a function of beacon SNR and re-synchronization interval for seven transmitters synchronized via the two-way carrier synchronization protocol.

from the noisy observations. After synchronization is complete in this example, clock offset standard deviations better than 100 ps and frequency offset standard deviations better than 0.4 parts-per-billion (PPB) can be obtained when the beacon signal-to-noise ratio (SNR) is greater than 5 dB and the re-synchronization interval $T \leq 500$ ms.

Distributed wireless carrier synchronization techniques achieve high accuracy by exploiting the timing information contained in the phase and frequency of the bandpass beacons exchanged during synchronization. While distributed wireless carrier synchronization can, in principle, achieve much more precise clock and frequency offset than PTP and GPS, the use of unmodulated beacons can lead to periodic ambiguities in the phase estimates. Hence, distributed wireless carrier synchronization techniques may be used in conjunction with conventional lower-precision synchronization techniques for to provide appropriate synchronization at different timescales. For example, GPS can be used with two-way carrier synchronization to provide symbol synchronization and also to stabilize the frequencies of the carriers at the BSs.

8.1.5 Summary

In this section, we have introduced the concept of synchronization and described several approaches to the problem of synchronizing nodes in a coordinated communication network. These techniques can be used separately or in conjunction to facilitate the establishment of a common notion of both frequency and time to a desired level of accuracy. The following sections will now analyze the effect of residual time and frequency errors on the performance of coordinated communication networks.