Chapter 10

Conclusions

Contents

10.1 General observations

The subjects of the Chapters in this Handbook are interdependent. Describing one topic without calling upon information from another is difficult. While the author of each chapter does not often refer explicitly to other chapters, the reader will no doubt find a need to do so. A key conclusion then is that the pursuit of work in one topic will probably require development of understanding of others. Consider the following examples of interdependence.

Advanced clocks and oscillators cannot be readily understood without an understanding of earlier clocks and oscillators, and work on either current or advanced devices surely requires an understanding of both the statistical basis for characterisation and methods of measurement.

To appreciate time and frequency applications, an understanding of several topics might be needed. These include the characteristics of clocks and oscillators, the construction of reliable time scales, the dissemination of time and frequency signals, the characterisation of signals transferred through such dissemination systems, and practical problems associated with these activities. For example, the performance of received signals often depends on the characteristics of the local broadcast source, the time scale to which that source is steered, and the dissemination process.

10.2 Clocks and oscillators

The range of clocks and oscillators available for applications is excellent. Commercially available clocks and oscillators include, for example; quartz, rubidium, cæsium, and hydrogen devices. This is the approximate ascending order of cost of these timing devices. Their basic principles of operation (outlined in Chapter 1) and their performances (outlined in Chapter 2) vary considerably. Specific application requirements (frequency stability, drift, sensitivity to environment, size, cost, weight, power requirements, etc.) might suggest selection of any one of these devices. Beyond these current commercial solutions to application requirements lie research advances (outlined in Chapter 9) that suggest future availability of yet more advanced devices. The paragraphs below summarise some key conclusions in this area.

Quartz oscillators form the very foundation of time and frequency technology. Passive atomic standards are critically dependent on quartz oscillators. In fact, the attack time of the servo systems in these standards is long. This means that the short term stability of passive atomic standards is essentially that of the quartz oscillator used to control the microwave oscillator that probes the atomic resonance. The number of quartz oscillators used in even very high-technology applications is enormous. They clearly dominate atomic standards in numbers by several orders of magnitude. While their long-term stability and drift are inferior to those of atomic standards, quartz devices remain desirable solutions because of their low cost, low weight, small size, high reliability, and low power requirements.

Rubidium frequency standards occupy the next level up in performance. Their stability is usually better than that of quartz oscillators for times longer than 10^4 s, but is not nearly as good as that of beam standards. They are generally less sensitive to environmental changes than quartz. While the packaged rubidium standard is typically larger than a packaged quartz oscillator, it is much smaller than currently available cæsium beam standards. Rubidium frequency standards are finding ever more applications, particularly in telecommunications systems, and have performed well as on-board clocks on the GPS satellites.

Cæsium beam frequency standards occupy a unique niche, because they rely on the resonance that serves as the definition of the second. Consequently, cæsium standards are the only choice for national primary frequency standards. The primary frequency standard differs from the typical field (commercial) standard in that systematic errors are carefully evaluated to arrive at the best possible realisation of the second. In practice, many systematic effects are under good control in commercial cæsium beam standards, so they can provide accuracy, but at a level below that of the primary standards. Because systematic effects are under better control, cæsium beam standards typically exhibit less drift than other standards, and the user can have more confidence that the output frequency corresponds closely to that stated by the manufacturer. Thus, cæsium beam standards are usually the first choice for systems where accurate frequency must be autonomously kept for long periods. Cæsium beam standards play a major role in the ground-control stations of GPS and are also carried (along with rubidium standards) by the GPS satellites. At present (1995), most of the clocks providing time and frequency from the GPS satellites are cæsium.

While priced much higher than other frequency standards, hydrogen masers deliver far superior stability in the short and intermediate term. They yield only to cæsium in the long term, and even then not always. If the resonance of the microwave cavity is servo controlled, a hydrogen maser can be more stable than a commercial cæsium standard for times of the order of one year. The long-term-stability problem is a result of the difficulty in controlling systematic effects. One of the most troubling of these is the shift associated with collisions of the hydrogen atoms with the walls of the storage bulb. Another is the pulling of the hydrogen-frequency resonant cavity, causing an error in the output frequency. The short-term stability of hydrogen masers has been extremely useful in applications such as very long-baseline interferometry (VLBI) where the time tagging of closely spaced observations is especially critical. Masers are also useful for characterising other high-performance frequency standards, since the time needed to make a given measurement becomes smaller as noise in the reference device decreases. The typical hydrogen maser is active, that is, it oscillates spontaneously, but passive hydrogen masers have also been built. Passive masers are intermediate in performance between active hydrogen masers and cæsium standards.

Research is under way on all these standards. Quartz devices have been improving gradually with time, and gradual improvement in the future is likely to follow improvement in understanding of materials and noise mechanisms, and development of clever methods for controlling environmental sensitivities. On the other hand, research in atomic physics has provided new tools that promise improvements of several orders of magnitude in performance of atomic standards. Lasers can now be used to control the atomic states and motions of atoms, dramatically reducing limitations related to atomic motion (Doppler shifts) and observation time. The ion frequency standard and the cæsium fountain top the list of advanced passive standards. Rubidium standards will likely benefit from laser pumping techniques that could replace discharge-lamp pumping, since performance limitations are ascribed to the traditional pumping method. Hydrogen masers should benefit from research on cryogenic wall coatings that promise substantial improvement in their already impressive short-term stability. Finally, mention should be made of research on the cooled sapphire resonator and the superconducting microwave resonator, both of which can exhibit exceedingly high Q, consequently leading to short-term stability well beyond that of the hydrogen maser.

10.3 Measurement methods and characterisation

Measurement methods in this field are dictated by the nature of the noise encountered in clocks, oscillators, and signal transfer systems. The realisation that lower-frequency noise is not usually white, but varies as the reciprocal of some higher power of the frequency, demands non-standard statistical treatment of the noise. Since timekeeping is a long-term activity requiring consideration of very long-term (low-frequency) behaviour, the statistical treatment of these non-white noise processes is particularly important. The random noise is clearly important, but systematic effects must also be understood, characterised, and controlled. Over the last several decades these subjects (well described in Chapters 3, 4, and 5) have evolved substantially.

Both time-domain and frequency-domain characterisation methods are important. Time-domain measures (described in Chapter 3) are especially useful for characterising long-term processes and frequency-domain measures are more useful for characterising shorter-term (higher-frequency) behaviour. The two-sample (Allan) variance and variations of it have satisfactorily replaced the standard variance, which cannot be used because it diverges for some non-white noise. Two-sample variances depend on the averaging time, so this characterisation approach involves a graph, a $\sigma - \tau$ plot, rather than just a one-number variance. The variance measures lend themselves well to the technology, since repeated measurements with a counter can be used to acquire data. Averaging time τ is then adjusted in the software processing of the data. The τ dependence is related to the Fourier frequencies present, hence a $\sigma - \tau$ plot can give a quick indication of the frequency components present in the data.

Shorter-term noise is more typically described using a spectral purity measure such as power spectral density. Phase-modulation (PM) noise is typically of greatest concern, but amplitude-modulation (AM) noise should not be completely overlooked. Chapters 3 and 4 provide descriptions of both the characterisation concepts and the methods for making the physical measurements. Direct measurements using a spectrum analyser can be useful, but the very high performances of many systems require higher resolution, which is usually achieved using a heterodyne method to bring the noise down from the higher frequency to base band.

Clocks and oscillators are subject to systematic effects induced by variations in such standard environmental parameters as temperature, humidity, barometric pressure, and magnetic field (see Chapter 5). Acceleration, vibration, shock, and ageing are also of concern. The physical origins of the responses of the various types of standards to changes in these environmental parameters are obviously different. An understanding of the relative magnitudes of these effects in the various devices can be of help in selecting a clock/oscillator for a particular application or in taking necessary precautions to reduce the effects of variations in the environment. It is also useful to understand the origins of these systematic effects. This requires study of their physical principles (Chapters 1 and 2).

Characterisation of clocks and oscillators is multifaceted, and errors are often made. A detailed understanding of the statistical measures and the measurement methods is essential to good metrology. Errors can result from mistakes in zero-crossing detection, truncation of data, improper accounting for bandwidth effects, and improper consideration of counter dead time. Measurements can be mathematically converted between the time and frequency domains, but considerable caution must be exercised in doing so, especially when going from the time to the frequency domain. Misunderstanding of characterisation concepts has led to so many problems that several organisations have developed standard nomenclature for terms and standard methodologies for performing measurements.

10.4 Time scales, coordination, and dissemination

For many applications the usefulness of the frequency output of an oscillator or the time output of a clock is not realised unless that signal is reliably maintained, compared with other oscillators/clocks and distributed to other locations (network-type applications). Many applications require continuous availability of time and/or frequency traceable to a central source. The questions that naturally arise involve methods for reliably maintaining the sources and optimum means for distributing signals from those sources.

In recent years, the ensembling of clocks with a proper algorithm to form a time scale has emerged as an important concept (Chapter 6). A properly implemented time scale can provide higher reliability and performance beyond that of any of the clocks contributing to the scale, and the means for assessing the performance of component clocks and even the time scale itself. Ensemble time scales rely on many earlier developments including methods for clock characterisation, low-noise methods for reading the outputs of clocks, and algorithms for optimally combining the outputs of an ensemble of clocks.

To maintain an autonomous time scale with an output related to UTC or any other time scale, co-ordinating the time-scale with the reference through time transfer is critical. Satellite methods have dramatically improved Coordination to the point where the best standards and time scales can be compared at their full accuracy, although this might require a number of days of averaging for the highest performance systems. Thus, the highest available accuracy can now be transferred anywhere. If an application requires continuous local availability of time, maintaining a high-stability clock system (preferably an ensemble) is best. The system can be steered with a long time constant to an external source. If the external signal is lost for a period the local system can adequately bridge the gap.

Most requirements for time and frequency signals involve accuracies well below those of the major reference time scales and simpler, more cost-effective methods of delivery of these signals (as described in Chapter 2) can be used. Several different signal delivery systems have been used over the years. These include telephone, HF radio, LF radio, and satellite signals. These are typically operated by national standards laboratories. In addition, other stable signals, not usually considered as standards, are often used as frequency references. A standard's laboratory can monitor such signals and provide correction data in the form of offsets from their own standard. The user applies these corrections to the received signal to achieve an accuracy that is traceable to the reference standard. Television and LORAN-C broadcasts have often been used in this way. Such methods allow a national laboratory to provide excellent service without a large investment in broadcast equipment.

The need for time stamping of business transactions and technical data by distributed computer systems has been growing rapidly. Current telephone and network time services are meeting some of these needs, but it is likely that timing signals will eventually be provided for such applications by telephone companies through fibre and cable systems.

10.5 Realities

Applications typically drive the development of a technology, and this is certainly the case (as discussed in Chapter 7) for time and frequency technology. Applications range from the very demanding scientific experiments to clocks and oscillators for a variety of simpler consumer products. Clocks and oscillators control the activities of many systems, that is, they regulate the rate at which events take place. As our societies have become more dependent on electronics, timing of activities has become progressively more exacting.

Frequency is our most accurately realised physical quantity. Because frequency measurements of modest accuracy are easy to make, other quantities are often transduced (converted) to frequency. For example, length is now defined as the distance travelled by light in a fixed time. Also, the Josephson volt standard converts frequency to voltage. These are examples of conversions of major measurement units, but many other types of transducers convert other quantities (for example, pressure and vibration) to frequencies. Thus, the value of frequency standards extends well outside the traditional areas of timing.

The close relationship between length and time has been exploited in navigation systems; the Global Positioning System (GPS) is particularly noteworthy. An accurate time interval translates directly to an accurate distance. In fact, the development of atomic frequency standards has been critical to the development of these navigation systems. The impact of these developments has been large. GPS, for example, is used not only for navigation by ships and aircraft, but for construction surveying, and in geophysical research. The ability to determine position accurately will no doubt be applied to many other activities generating a wide range of consumer products for positioning.

Telecommunications is another key area of application. Timing requirements are escalating as systems move to ever higher data rates. With the explosion in information transfer and processing, time and frequency technology is playing a larger role in telecommunications. The capabilities of atomic frequency standards have long been required in this industry, and wide distribution of high-level timing is critical to the successful operation of newer systems.

While the list of applications goes on, it might be well to conclude here by noting that this technology is also important to scientific research. Many fundamental physical theories yield predictions of clock behaviour, and clocks thus play a major role in tests of these theories. In fact, scientific applications have often served as the driving force for the development of better clocks.

This Chapter is concluded with a few words of caution. This is a complex technology involving a range of disciplines. These include, for example, atomic physics, materials science, electronics, measurement science, statistical analysis and satellite and terrestrial radiocommunications. Stability, reliability, and accuracy have come to have special meanings and importance in the field. These must all be solidly understood. Practitioners of the technology have accumulated experience suggesting that certain types of problems occur repeatedly. Some of these problems are discussed in Chapter 8. If you are about to embark on a major project in the field, seek the experience of others. Chapter 8 might provide a good starting point.