Advances in Metrology and Its Role in Quality Improvement and Global Trade
B.S. Mathur, V.N. Ojha and P.C. Kothari (Eds)

Time and Frequency Metrology

Donald B. Sullivan and Judah Levine
Time and Frequency Division, National Institute of Standards and Technology
Boulder, CO 80303, USA

ABSTRACT

This paper presents an overview of some of the trends which are shaping the directions of development in time and frequency metrology. The paper focusses on the characteristics of frequency sources, methods for comparing separated clocks/oscillators, and techniques for distributing time and frequency signals.

1. INTRODUCTION

Recent progress in time and frequency metrology suggests that developments taking place today are more far-reaching than any of those of the last two decades. The development of laser methods for manipulating the states and motions of atoms will provide dramatic advances in the performance of atomic standards. Advances in satellite time transfer now provide for transfer accuracy and stability well beyond the performance of most atomic frequency standards. Further progress in time transfer should support the application of even the most advanced frequency standards. In parallel with this activity, there is greatly renewed interest in substantially improving the synchronization of telecommunications networks, navigation networks, and electrical power networks. This paper provides a general look at these trends and, when practical, suggests where they are leading.

2. BACKGROUND

a. Accuracy and Stability

We first emphasize the difference between accuracy and stability. Consider the performance of a cesium-beam frequency standard. The accuracy of the standard describes its ability to generate a frequency with known systematic uncertainties (frequency shifts) relative to the ideal (the model). An accuracy statement involves an upper and lower limit for deviations of the standard from the model. In simple terms, the frequency stability of the standard is a measure of its ability to stay within specific frequency limits for some sampling time τ.

The evaluation of the accuracy of a standard is based on a physical model of the standard. The accuracy statement involves a proper combination of all of the errors derived from independent measurements and the theoretical predictions of the model. Such evaluations are inevitably checked through comparisons among the independently developed primary standards of the world. If a standard is highly accurate, it obviously has very good long-term stability. Many of the world's standards laboratories take the unproven, but intuitively appealing, position that the best approach to long-term stability is to improve accuracy. The basis for this position is the idea that long-term variations in output are caused by variations in the systematic offsets. By reducing and controlling the offsets, we thus improve both accuracy and long-term stability. Clearly, the accuracy of a standard can be no better than its long-term stability.

Outside the standards laboratory (in practical situations) stability is often the key consideration. For example, if several nodes in a telecommunications system must be properly timed for synchronous communication, it matters little whether the time delivered to the nodes is accurate. All that really matters is that all nodes measure the same time. If, however, the network is large (many nodes) and synchronization is acquired from several sources, then it may be necessary to require accuracy as well.

Kartaschoff and Barnes [1] present a broader discussion of accuracy and stability. The measures of stability used in this field have matured to the point where they are now the subject of an IEEE standard [2,3]. The precursor to this standard is a highly referenced paper by Barnes et al. [4]. In the historical development of these measures, time-domain methods have dominated.

b. Frequency Standards

The major advances in timekeeping in this century have been the developments of the quartz-crystal oscillator and the atomic clock. Quartz oscillators are used in almost all timing systems. The number of quartz oscillators produced annually is estimated to be on the order of 10^9. Quartz-oscillator technology has been evolving steadily over a long period. The likelihood is that steady improvement will continue as better crystals and cleverer compensation schemes are developed.

In order of increasing cost, the four main classes of currently used standards are quartz oscillators, rubidium frequency standards, cesium-beam frequency standards, and hydrogen masers. In general, this is the same as the ascending order of performance, except that the cesium standards provide the best accuracy and long-term stability while hydrogen masers provide the best short-to-medium-term stability. The number of atomic standards produced annually is measured in thousands of units.

The atomic clocks of today share two common characteristics. (1) The atoms in all of them move at thermal velocities commensurate with room or higher temperatures. This means that Doppler shifts (and corrections for them) are extremely important. (2) The required difference in the population of atoms in the excited and ground states in clocks is achieved through a pumping or selection process which itself adds complications. For example, the Stern-Gerlach magnets used in cesium standards introduce a transverse velocity dispersion which, through the cavity phase shift, produces a frequency error. In another example, the discharge-lamp pumping of rubidium standards involves excess, broad-spectrum light which complicates the pumping process. The aging of the discharge lamp contributes to drift and instability.

We are now on the threshold of a revolution in the design of atomic frequency standards. At the heart of this revolution lies laser control of the motions and atomic states of ions and atoms. General aspects of this revolution are described in section 3.

c. Time Transfer Systems

Ten years ago, the methods for time comparison of widely separated clocks involved uncertainties which were greater than the performance of the best primary standards. In this environment, the development of better standards was largely an academic exercise. Signals could not be reliably transferred to other locations, and accuracy claims could not be checked by comparison with other standards. The development of accurate satellite time transfer systems, which has taken place over the past two decades, has changed this. Common-view time transfer using the Global Positioning System (GPS) satellites [5] has proven to be extremely accurate, offering global coverage with an uncertainty on the order of 10 ns. The two-way method for time transfer [6,7], which uses standard satellite communication channels, promises even higher performance. These systems put the accuracy and stability of time transfer (or time comparison) ahead of time generation. The development of better standards is thus no longer academic.

d. Network Synchronization

One of the key practical challenges to time and frequency technology is the synchronization of nodes in major networks such as computer networks, telecommunications networks, electrical power networks, and navigation networks. In the transfer of information in a communications system, synchronous operation provides for higher throughput [8]. Synchronization in electrical power networks supports fault location, event recording, and control of system stability [9]. Better synchronization in navigation networks obviously results in more accurate navigation.

Systems for network synchronization can be divided into two broad categories: (1) peer organizations in which groups of nominally identical nodes exchange time data so as to establish a single self-consistent time scale and (2) stratified, client-server arrangements in which the time of each node is obtained from a relatively small number of sources or possibly a single primary source. Using time servers reduces the time distribution problem to the determination of the errors introduced by the individual delays between the source and the network nodes. The success of the method will depend on how well these errors can be determined. Although groups of peers must also determine the delays between the nodes, the increased symmetry of this configuration may simplify the problem somewhat since the time of a single node is generally compared to the time of several other nodes connected by different paths with different delays. At least in principle, it is possible to perform a dynamic least-squares adjustment which minimizes the average errors over all the nodes. In either configuration, the reliability and short-term stability of the system are substantially improved if the central timing signal is used to steer a local clock at the node rather than to directly control the operations at the node. If the central synchronization signal is interrupted, the local clock can carry the system at that node for some period before synchronization is completely lost. Furthermore, the higher frequency noise inherent in time transfer systems substantially compromises the short-term stability of the as-received signal. Virtually all network synchronization relies on this approach (the control of a local clock by the synchronization signal). The quality of the local clock and the sophistication of the steering are choices dictated by the reliability required in the application.

Where reliability of synchronization is paramount, a second (or even third) independent distribution system (with or without a second reference clock) can be used. With two synchronization signals delivered to each node, the loss of one distribution system (or central clock) simply forces each node to rely on the alternate source.

3. TIME GENERATION

Current studies at a number of laboratories are demonstrating principles which, when applied to standards and oscillators, will certainly result in improved performance. For example, Doppler shifts can now be minimized by reducing velocities of beams of atoms and the thermal motions of trapped ions. Newer methods for optical state selection and detection eliminate some of the negative side effects of current methods. Details on these methods can be found in a number of papers. Itano and Ramsey [10] present a brief history of the development of atomic standards and project the future development of these standards, while Itano [11] and Rolston and Phillips [12] describe progress toward cooled-ion and cooled-atom standards. Lewis [13] presents a comprehensive description of the various types of atomic standards and their characteristics. In the sections below, we present brief discussions of the various types of clocks and oscillators, indicating the status and directions of work.

a. Cesium-Beam Standards

Primary cesium-beam standards are approaching certain practical limitations. For example, in a typical primary standard the correction for the second-order Doppler shift is about 2×10^{-13}. For a system with a design accuracy of 1×10^{-13}, this correction is not too difficult to make, but at 1×10^{-14} it becomes a substantial problem. The narrowest linewidths for the cesium clock transition (~9.2 GHz) are typically tens of hertz, so at an accuracy of 1×10^{-14} the clock servo system must find line center with an accuracy approaching 1×10^{-6}. These are but a few of the reasons we cannot expect to see the performance of these standards go much beyond 1×10^{-14}. Drullinger [14,15], Ohshima et al. [16], and de Clercq et al. [17] describe thermal beam systems which use optical state selection and detection rather than the conventional magnet selection. This approach, as well as other variations on the conventional technology, should allow achievement of the 10^{-14} accuracy, but they do not really avoid the problems noted above.

The real solution to these limitations lies in slowing the atoms so that the Doppler shift is reduced and longer observation times can yield narrower resonance linewidths. Itano [11] describes a most direct solution, the trapping and cooling of positive ions. This concept has matured to the point where high-performance prototype standards have been demonstrated. Rolston and Phillips [12] discuss an approach involving the slowing of neutral atoms. The lack of a suitable trap for neutral atoms limits the achievable linewidth. One proposal for a cooled-atom standard involves a fountain where slowed atoms are lofted vertically and interrogated as they rise and then fall under the influence of gravity. The potential signal-to-noise ratio with slowed neutral-atom standards is better than with trapped ions. Furthermore, the definition of the second is now based on neutral cesium, so primary standards based on slowed cesium atoms would continue to be favored unless other standards prove to be greatly superior.

The two areas in which new concepts will likely affect practical cesium standards (field standards) in the next ten years involve optical state selection/detection and closed-cell standards in which cesium atoms are cooled sufficiently to reduce transition linewidth and Doppler shift. A key advantage of optical state selection and detection is that they can support the use of all atoms in the atomic beam. In conventional magnetic selection, 15/16 of the beam is in the wrong atomic state and has to be discarded. The more efficient use of the beam atoms means that signal-to-noise ratio can be greatly increased (at the same beam flux), or the lifetime of the standard can be increased by operating at lower oven flux (while maintaining respectable signal-to-noise performance). The cesium-cell concept developed by Monroe et al. [18] uses slowed cesium atoms contained in a closed envelope similar to that used in rubidium standards (see the discussion of rubidium standards by Lewis [13]). The cell approach is likely to result in a simpler overall system (once the laser diodes are sufficiently simplified) of very good medium-term stability.

b. Stored-Ion Standards

The most readily understood advantages of using trapped ions as frequency standards are that (1) the first-order Doppler shift is eliminated because the ions remain fixed in position and (2) the transition linewidth is dramatically reduced by long observation times. Furthermore, trapped ions can be readily laser cooled to minimize the effect of the second-order Doppler shift. The use of radiation pressure to cool trapped ions was first demonstrated in 1978 by Wineland et al. [19] and Neuhauser et al. [20]. The cooling of neutral atoms, first demonstrated in 1981 [21,22], also uses radiation pressure. In these ion and neutral-atom experiments the fundamental cooling concept is the same, but the implementations are different. The ion-storage technology is clearly more mature, and demonstrations of high accuracy and high stability performance already push the performance of conventional cesium standards.

Bollinger et al. [23] demonstrated the first high-accuracy prototype ion standard using Be^+ ions. Cutler et al. [24] have demonstrated a highly stable ion standard (not laser cooled) which outperforms cesium in the medium term. Prestage et al. [25] have further developed this type of standard. But the real promise for ion standards still lies in the future. The NIST group [26], in studies of an optical-frequency transition in single trapped Hg^+ ions, anticipate that systematic frequency shifts can be determined to 1×10^{18}. The linewidth of this optical transition is 2 Hz, providing an inherent Q factor of 10^{15}. Bergquist et al. [27] have locked a laser to this transition and achieved a linewidth of 80 Hz limited by the linewidth of the laser. This represents the narrowest laser linewidth and the highest-Q atomic or molecular transition ever observed. Optical transitions are appealing for frequency standards because they offer the highest Q factors, but there is a fundamental problem in working in the optical region. This is the difficulty of precisely relating the output frequency to a frequency in the microwave frequency region where the performance can be applied to practical electronic metrology. Frequency multiplication (or division) between the microwave and optical regions is extremely difficult. It must be dramatically simplified before optical frequency standards are widely accepted. The Hg^+ ion also has a clock transition at 40.5 GHz. This can be used in standards that do not look too different (electronically) from conventional rubidium standards. Using laser cooling to narrow the linewidth, a clock based on this transition should have an uncertainty of $<10^{-16}$.

One disadvantage of ion standards is that the most accurate single-ion systems exhibit poor signal-to-noise ratios. Prestage et al. [25] have shown that a modification of the usual trap to one of linear geometry allows trapping of a larger number of ions without a substantial decrease in the potential accuracy. The larger number of ions provides for a higher output signal and a higher signal-to-noise ratio. Preliminary experiments suggest that ion standards based on a linear trap might even challenge the active hydrogen masers (see below) in short-term stability.

c. *Other Atomic Standards*

Two other classes of atomic standards, hydrogen masers and rubidium standards, play a significant role in science and technology. Hydrogen masers are the common choice where very high short-term stability is required, and rubidium standards now provide a cost-effective performance which is better than that of quartz oscillators, although below that of cesium. (Superconducting-cavity-stabilized oscillators [28] actually show the best short-term stability, but they are not yet in common use.)

There are two types of hydrogen masers. Active hydrogen masers are distinguished by the fact that the atomic gain is sufficient to produce spontaneous oscillation. Passive masers use an external oscillator to probe the hydrogen resonance. Active masers are critical to very-long-baseline interferometry (VLBI), where observations of radio-emitting stars must be time tagged with very high short-term stability. Passive hydrogen masers are smaller (comparable to the size of a cesium standard) and less expensive than active masers. Their niche is intermediate-term stability, which is generally better than that of cesium standards.

Much of the research surrounding present maser development involves the wall coatings of the bulbs which contain the hydrogen within the microwave cavity. A key characteristic of masers is the very long interrogation time of the hydrogen transition which results from the fact that individual hydrogen atoms can suffer many collisions with the walls in the bulb without disturbing the atomic state in a major way. There is, nevertheless, an energy (frequency) shift associated with the wall interaction. The stability of this shift has long been a source of study and discussion. Wall coating, often more an art than a science, has played a key role in maser performance. Future advances in maser performance are critically dependent upon improvements in these coatings. Masers developed recently in the Soviet Union [29] have performances which suggest a substantial improvement in wall coating. Another approach to the wall problem has been the development of masers which use liquid ^3He as the wall coating [30,31]. Preliminary experiments with these cryogenic masers look promising. Yet the technology is still very difficult, and it will probably be many years before such masers are ready for practical applications.

Conventional rubidium standards, optically pumped by special discharge lamps, use a passive buffer gas to slow the diffusion of rubidium atoms. This increases the atomic lifetime in the interrogation region resulting in a narrower linewidth. Because these standards perform better than quartz oscillators and sell for only a fraction of the price of cesium standards, they fill an important gap in the technology. However, they suffer aging effects which give rise to substantial drift, and the excess spectrum of light from the pump lamp reduces the signal-to-noise ratio.

Recent research [32,33] suggests that some of these problems might be overcome through replacement of the discharge lamp with a suitably controlled diode laser operating at the rubidium transition frequency. This appears to minimize (or even eliminate) both the lamp-aging problem and the excess light which affects the signal-to-noise ratio. Some people suggest that the rubidium standard might some day achieve a short-term performance challenging that of today's active hydrogen masers.

4. TIME DISTRIBUTION

The transfer medium plays a critical role in time distribution. For example, short wave (HF) distribution of time signals is cost effective, but with the unpredictable bouncing of the signals between the earth and ionosphere, it is an inappropriate method for high-accuracy dissemination. In another example, the transmission delay for satellite-distributed time signals depends strongly on the carrier frequency. Ionospheric variations in delay can be tens of nanoseconds at the 1 GHz GPS frequencies. These variations are lower by a factor of 100 in the direct-television-broadcast band (satellite band at 12-14 GHz). Clearly, the selection of the linking medium is an important consideration in network synchronization.

a. General Concepts

The general concepts for time transfer are easily sorted into three categories, one-way, common-view, and two-way time transfer. The problem in every case is how to best deal with the delay introduced by the transmission medium, be it a radio-wave path, an optical fiber, or a coaxial cable.

The one-way delivery of a time signal from a source to a user is typified by short-wave radio broadcasts such as those emanating from WWV in Colorado. The delays in such broadcasts, which can be tens of milliseconds, are highly variable because the transmission path often involves multiple reflections between the ionosphere and the surface of the earth. The accuracy of one-way transfer can be improved by characterizing the delay of the medium, but in the case of short-wave broadcasts, this characterization involves a large uncertainty. On the other hand, one-way signal delivery through a coaxial cable or from a satellite involves a much more predictable path, and delay corrections can be stable. Cable delays are affected, for example, by temperature and stress on the cable, while satellite delays are affected by dispersion in the ionosphere and troposphere and, of course, changes in position of the satellite. In general, one-way systems are limited by the difficulty in obtaining a complete characterization of delay variations.

The common-view method of time transfer can be used to advantage where the path from some reference source to each of two receivers involves a path delay with common characteristics. For nearly simultaneous observations of the common source, the two sites record time differences which are $A - (R + \tau_{ra})$ and $B - (R + \tau_{rb})$, where A, B, and R are the readings of clocks at Site A, Site B, and the Reference. τ_{ra} and τ_{rb} are the respective delays between the Reference and Sites A and B. If the two sites then share their readings and take the difference between them, the difference, $A - B - (\tau_{ra} - \tau_{rb})$, does not contain the common reference, but leaves only the differential delay, $(\tau_{ra} - \tau_{rb})$, as a correction. Clearly, the more common the characteristics of the path, the better will be the time transfer. Furthermore, the accuracy of the time transfer does not depend (to first order) on the quality of performance of the common reference clock. Allan et al. [34] have used this technique with the Global Positioning System (GPS) satellite clocks as references. While time directly received from these space-borne clocks is no more accurate than about 200 ns, they achieve a time transfer accuracy on the order of ±10 ns. Careful error correction and averaging of at least a day's observations are needed to achieve this performance.

The common-view method can be extended to more than two receivers, and there may be advantages to such an arrangement. The data from each receiver can be combined with many others to yield simultaneous common-view estimates of the time differences of all members of the ensemble, and a least-squares adjustment can then be performed to minimize the error at each node. The power of this procedure will depend on the correlations among the various differential delays.

The two-way method for time transfer provides the best opportunity to accurately determine the transmission delay between separated clocks. When sites A and B simultaneously exchange time signals through the same medium and compare their received readings with their own clocks, they each record the respective differences, $A - (B + \tau_{ba})$ and $B - (A + \tau_{ab})$, where τ_{ba} is the transmission delay from B to A and τ_{ab} is the reverse delay. Taking the difference between these two sets of readings we obtain $2(A - B) + (\tau_{ba} - \tau_{ab})$. Now, if the transmission path is fully reciprocal, that is, if $\tau_{ba} = \tau_{ab}$, then the difference, A - B, is known perfectly, and the value of the transmission delay is also determined. The two-way method should provide for nearly real-time synchronization of clocks at nanosecond or lower uncertainty.

The common-view and two-way methods impose different requirements on the transmission medium. The common-view method depends on the relatively high correlation between the paths from the transmitter to the two receivers. The paths may include arbitrary delays and need not be reciprocal. The two-way method, on the other hand, places a premium on reciprocity and does not depend on the correlation between the delays in different portions of the path. Although there are areas of overlap, each technique also has a domain in which it can provide better performance. The common-view method, for example, is likely to be more robust in computer networks, which have large transmission delays and poor reciprocity. The two-way method is probably the method to choose when signals follow cable or line-of-sight paths.

b. Applications of the Time Transfer Concepts

In the discussion above we used several examples to illustrate the different time transfer methods. But the concepts are general, and each can be applied to a variety of transmission schemes. The discussion in this section is not meant to be comprehensive, but rather to indicate a few of the real and potential applications of the concepts.

Jackson and Douglas [35] and Levine et al. [36] describe two-way methods for synchronizing remote clocks through the telephone system. They both use the following variation of the two-way idea. Rather than accomplish the two-way exchange through simultaneous transmissions, these systems send a time signal from site A to site B where it is echoed (reflected) back to site A. Site A then has a measurement of the round-trip delay which, assuming reciprocity in the lines (an assumption uncertain within 1 ms for many systems), is divided by 2 in order to provide a direct measure of the delay. The main difference between the methods described in the two publications above is that the first takes care of the correction at each receiver and the second makes the corrections at the transmitting end.

Using a similar scheme in cable-connected systems, two-way time transfer can push the stability and accuracy of synchronization into the picosecond regime, even for systems with very long cables. Where the demand for accuracy/stability of distribution of time within a site are high, the use of this method along with optical fiber connections [37] is especially appealing since the fiber transmission provides excellent immunity to noise introduced by ground-isolation problems.

Navigation or ranging systems which are based upon time of flight of radio or optical signals can often be used for time transfer. For ranging applications these systems assume that the time references are known and solve for relative position(s) of base stations. For time transfer, the known locations of the base stations are given and the relative time at the stations is then determined. This is the basis for the GPS Common-View Method of time transfer described earlier [5]. The common-view technique of very-long-baseline interferometry (VLBI), used to obtain high-resolution images of distant radio stars or relative position(s) of observation sites, provides another means for highly stable time transfer [38]. In this technique, signals received at two sites from a set of common radio stars are cross correlated (in a computer) in the same fashion that optical signals are combined in an optical interferometer. A key product of this computation is a very precise measurement of the relative phase of arrival of signals at the two sites. This can be used to provide an extremely stable measurement of the relative stability of clocks operating at the two sites. With the very high cost of radio telescopes and peripheral equipment and the need for extensive computer processing of data, this is an expensive approach, and it is generally used only as a research tool.

For satellite time transfer, the one-way method remains attractive because of the simplicity of equipment at the receiving site [39]. Two-way time transfer requires broadcasting to a satellite from each station, a process which requires more equipment and special licensing. The common-view methods call for exchange of information between sites as well as substantial averaging to obtain good accuracy and stability.

c. *Practical Synchronization Limits*

Given adequate clocks, the synchronization of two nodes or a network is likely to be limited by the transmission medium and/or the time-transfer method. Often it is possible to gain auxiliary information which might allow for a refinement of the estimate of time delay. For example, knowledge of the temperature of a coaxial cable along with an understanding of the delay-temperature relationship can provide for an improvement in one-way synchronization through coaxial cables. In another example, a model of ionospheric delays is currently used to improve upon GPS common-view time transfer. Furthermore, if variations in delay have suitable statistical properties, these can be used to place a confidence on a particular measurement or, with proper time constant, to better steer a remote clock to a central standard.

For the one-way method, practical limitations simply involve the variations in path delay and methods used to estimate them. Practical limitations of the two-way method are also readily understood. These limitations relate primarily to the assumption of reciprocity of the signal path through the medium and the transmitting and receiving equipment as well as the simultaneity of the exchange. The current approach to two-way time transfer through telecommunications satellites involves different frequencies for the up links and down links with the satellite. Since the delay is frequency dependent, the path is not fully reciprocal, although the errors introduced appear to be very small [40]. A general analysis of the limitations of common-view time transfer is more difficult for a number of reasons. The common-view cancellation of delay error is usually only partial, and a variety of additional methods are used to reduce the uncertainty in the determination of the remaining differential delay. These methods include straightforward averaging over many observations, cross correlation of the data obtained in observing a number of independent reference sources, and modeling or independent characterization of the transmission medium. Because of this complexity, the limitations of this transfer method involve more detailed consideration.

The synchronization limits imposed by the time transfer process are, in a certain sense, very general limits on time and frequency technology. The development of atomic clocks with performances beyond the time transfer limits makes sense only if such clocks are to be used in isolation. Applications for such isolated clocks are fewer, probably involving only scientific studies.

6. DISCUSSION/CONCLUSIONS

The development of methods for manipulating the states and motions of atoms and ions will certainly result in substantially improved atomic clocks. Since there do not appear to be any new principles to apply to quartz oscillators, we should expect performance improvement to be much less dramatic. There is a strong motivation for improving the performance of clocks only as long as time transfer techniques lead the performance of the clocks. Should clock development outpace transfer systems, then such high performance will be available only in the laboratory. Accuracy which cannot be transferred to another site is probably important only in specialized applications in science (for example, tests of relativity theory).

The one-way, common-view, and two-way time transfer concepts have been around for many years. The dramatic improvements in time transfer are more closely related to general technological developments which make the implementation of some of these concepts economically feasible. The areas of development which have had impact on time transfer include those in satellite communications, computers and microprocessors, and optical fibers. The focal problem in time transfer continues to be the characterization of the delay through the time transfer medium.

REFERENCES

* Work of the U.S. Government. Not subject to copyright.

[1] P. Kartaschoff and J.A. Barnes, Proc. IEEE, **60**, 493 (1972).
[2] IEEE Std. 1139-1988, 1988. (Available from IEEE Service Center, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331.)
[3] D. Allan, H. Hellwig, P. Kartaschoff, J. Vanier, J. Vig, G.M.R. Winkler, and N. Yannoni, 42nd Ann. Symp. on Frequency Control, IEEE Catalogue No. 88CH2588-2, 419 (1988).
[4] J.A. Barnes, A.R. Chi, L.S. Cutler, D.J. Healey, D.B. Leeson, T.E. McGunigal, J.A. Mullen, Jr., W.L. Smith, R.L. Sydnor, R.F.C. Vessot, and G.M.R. Winkler, IEEE Trans. Instrum. Meas., **IM-20**, 105 (1971).
[5] W. Lewandowski, G. Petit, and C. Thomas, IEEE Trans. Instrum. Meas., **42**, 474 (1993).
[6] D.W. Hanson, 43rd Ann. Symp. on Frequency Control, IEEE Catalogue No. 89CH2690-6, 174 (1989).
[7] Y. Saburi, M. Yamamoto, and K. Harada, IEEE Trans. Instrum. Meas., **IM-25**, 473 (1976).
[8] P. Kartaschoff, Proc. IEEE, **79**, 1019 (1991).
[9] R. Wilson, Proc. IEEE, **79**, 1009 (1991).
[10] W.M. Itano and N.F. Ramsey, Sci. Amer., **269**, 56 (1993).
[11] W.M. Itano, Proc. IEEE, **79**, 936 (1991).
[12] S.L. Rolston and W.D. Phillips, Proc. IEEE, **79**, 943 (1991).
[13] L.L. Lewis, Proc. IEEE, **79**, 927 (1991).
[14] R.E. Drullinger, Proc. IEEE, **74**, 140 (1986).
[15] R.E. Drullinger, D.J. Glaze, J.P. Lowe, and J.H. Shirley, IEEE Trans. Instrum. Meas., **44**, 120 (1995).
[16] S. Ohshima, Y. Nakadan, T. Ikegami, Y. Koga, R. Drullinger, and L. Hollberg, IEEE Trans. Instrum. Meas., **IM-38**, 533 (1989).
[17] E. de Clercq, A. Clairon, B. Dahmani, A. Gérard, and P. Aynié, Frequency Standards and Metrology, A. De Marchi, ed. (Springer-Verlag, New York, 1989) 120.
[18] C. Monroe, H. Robinson, and C. Wieman, Opt. Lett., **16**, 50 (1991).
[19] D.J. Wineland, R.E. Drullinger, and F.L. Walls, Phys. Rev. Lett., 40, 1639 (1978).
[20] W. Neuhauser, M. Hohenstatt, P. Tsohek, and H. Dehmelt, Phys. Rev. Lett., **41**, 233 (1978).
[21] S.V. Andreev, V.I. Balykin, V.S. Letokhov, and V.G. Minogin, JETP Lett., **34**, 442 (1981).
[22] W.D. Phillips and H. Metcalf, Phys. Rev. Lett., **48**, 596 (1982).
[23] J.J. Bollinger, D.J. Heinzen, W.M. Itano, S.L. Gilbert, and D.J. Wineland, IEEE Trans. Instrum. Meas., 40, 126 (1991).

[24] L.S. Cutler, R.P. Giffard, P.J. Wheeler, and G.M.R. Winkler, Proc. 41st Ann. Symp. on Frequency Control, IEEE Cat. No. CH2427-3, 12 (1987).

[25] J.D. Prestage, G.R. Janik, G.J. Dick, and L. Maleki, IEEE Trans. Instrum. Meas., 40, 132 (1991).

[26] D.J. Wineland, J.C. Bergquist, J.J. Bollinger, W.M. Itano, D.J. Heinzen, S.L. Gilbert, C.H. Manney, and M.G. Raizen, IEEE Trans. UFFC, 37, 515 (1990).

[27] J.C. Bergquist, F. Dietrich, W.M. Itano, and D.J. Wineland, Proc. 9th Int. Conf. on Laser Spectroscopy, Bretton Woods, N.H., M.S. Feld, J.E. Thomas, and A. Mooradian, eds. (Academic Press, San Diego, 1989) 274.

[28] G.J. Dick and R.T. Wang, IEEE Trans. Instrum. Meas., 40, 174 (1991).

[29] A.A. Belyaev, N.A. Demidov, B.A. Sakharov, V.Yu. Maksimov, M. Yu. Fedotov, and A.E. Yampol'skii, Measurement Techniques, 30, 767 (1987) and B.A. Gaigerov, L.P. Elkina, and S.B. Pushkin, Measurement Techniques, 25, 23 (1982).

[30] R.F.C. Vessot, E.M. Mattison, R.L. Walsworth, and I.F. Silvera, Frequency Standards and Metrology, A. De Marchi, ed. (Springer-Verlag, New York, 1989) 88.

[31] M.C. Hürlimann, W.N. Hardy, M.E. Hayden, and R.W. Cline, Frequency Standards and Metrology, A. De Marchi, ed. (Springer-Verlag, New York, 1989) 95.

[32] J.C. Camparo and R.P. Frueholz, J. Appl. Phys., 59, 3313 (1986).

[33] M. Hashimoto and M. Ohtsu, Frequency Standards and Metrology, A. De Marchi, ed. (Springer-Verlag, New York, 1989) 434.

[34] D.W. Allan, D.D. Davis, M. Weiss, A. Clements, B. Guinot, M. Granveaud, K. Dorenwendt, B. Fischer, P. Hetzel, S. Aoki, M.K. Fujimoto, L. Charron, and N. Ashby, IEEE Trans. Instrum. Meas., IM-34, 118 (1985).

[35] D. Jackson and R.J. Douglas, Proc. 18th Precise Time and Time Interval Meeting, 541 (1986). (Available from Time Service, U.S. Naval Observatory, 34th and Massachusetts Ave., N.W., Washington, D.C. 20392-5100.)

[36] J. Levine, M. Weiss, D.D. Davis, D.W. Allan, and D.B. Sullivan, J. Res. of NIST, 94, 311 (1989).

[37] L.E. Primas, R.T. Logan, Jr., and G.F. Lutes, 43rd Ann. Symp. on Frequency Control, IEEE Catalogue No. 89CH2690-6, 202 (1989).

[38] W.K. Klemperer, Proc. IEEE, 60, 602 (1972).

[39] A. Sen Gupta, A.K. Hanjura, and B.S. Mathur, Proc. IEEE, 79, 973 (1991).

[40] J. Jespersen, 43rd Ann. Symp. on Frequency Control, IEEE Catalogue No. 89CH2690-6, 186 (1989).