Metrologia **42** (2005) S90–S99 doi:10.1088/0026-1394/42/3/S10

Fifty years of commercial caesium clocks

Leonard S Cutler

Agilent Laboratories, 3500 Deer Creek Road, Palo Alto, CA 94304, USA

Received 14 February 2005
Published 7 June 2005
Online at stacks.iop.org/Met/42/S90

Abstract
The first commercial caesium-beam clock was the Atomichron, developed by National Company. The first unit was delivered in 1956. This paper first presents a brief history of commercial caesium-beam clocks. It then covers many of the sources of frequency error and frequency instability. Some details of the technical design of the caesium-tube and electronics are presented. Finally, we present some possible directions for future commercial caesium clocks.

1. Introduction

Commercial caesium clocks have been available now for almost 50 years. There have been a number of companies that have developed them for sale to both the military and commercial segments. Some of the companies are no longer in business and others no longer manufacture the clocks. This paper starts with a brief history of the commercial clocks from most of the manufacturers.

In this paper we use 'clock' and 'frequency standard' interchangeably since a frequency standard is an integral part of a clock. The output signal of a frequency standard is usually an ac voltage oscillating at some convenient frequency such as 5 MHz or 10 MHz. It may be sinusoidal or pulsed. To make a clock, the output signal of a frequency standard is applied to a clock mechanism, which integrates the standard's frequency with respect to time and scales the integral to provide time in the desired time-scale. The frequency standard controls the clock's rate. The time to which the clock is set is the constant of integration. Clearly, the most important part of the clock is its frequency standard.

Important performance aspects of commercial caesium-beam frequency standards include the accuracy, short- and long-term stability and environmental sensitivity. Factors affecting these will be addressed in this paper from a technical standpoint. Other aspects, such as cost, size, reliability and number, types and frequencies of the output signals, will not be covered.

Finally, some possible directions for future commercial caesium-standards will be presented. There are other atoms, ions or molecules that could be, and are, used but, because of the intent of this paper, the discussion will be limited to caesium.

2. History

Most of the available historical data are from National Company [1–3] and Hewlett-Packard Company/Agilent Technologies. The data from Hewlett-Packard/Agilent Technologies are most complete since the author of this paper has the most information to provide because of his involvement with their caesium clocks from the very beginning to the present. The history of other companies is scant and hard to get. What is presented here on the other companies is derived from personal contacts and websites. At present, Symmetricom and Agilent are the major suppliers of commercial caesium-standards based in the USA. Oscilloquartz is the major supplier based in Europe.

2.1. National Company

The first commercial caesium-beam frequency standard was built by the National Company [1, 2]. It was based on concepts of Professor Jerrold Zacharias at MIT. Richard Daly and Joseph Holloway developed the physics package, which included a compact new type of vacuum pump and an efficient caesium source. The physics package was called (and is today) the caesium beam-tube. It was about 2.1 m tall with a 0.1 m diameter at its largest point. It contained the microwave cavity, the caesium source, the vacuum pump, the hot-wire detector, the mass spectrometer and the electron multiplier. The deflection magnets were outside the tube. Ramsey-separated field excitation [4] was used with the two cavities fed by a waveguide with a tee-joint feed at the midpoint between the cavities. The caesium-beam travelled within this waveguide and thus was exposed to the microwave fields in what was supposed to be the field-free drift region. This did have some effect on the overall accuracy. The tube was placed vertically at the back of a rack cabinet on wheels. The electronics occupied the front of the cabinet. The whole cabinet was about 0.6 m deep by 0.6 m wide by 2.1 m tall. The electronics were all vacuum tube based. A klystron in a drift-cancelled scheme provided the final caesium microwave interrogation frequency multiplied from a 5 MHz quartz oscillator. The frequency standard was named Atomichron, and the first model number was NC1001. The first delivery occurred in 1956 and over the years through 1960, about 50 Atomichrons were sold.

The accuracy and long-term stability of the Atomichron were about $\pm 1 \times 10^{-10}$. The short-term stability is unknown. The reliability was claimed to be fairly good. A more compact version of the Atomichron was developed for the US Air Force. It is unknown how many of these compact units were built. The Atomichron assets, intellectual property and trade name were sold in 1969 to Frequency Electronics, Inc. (FEI) because of financial problems at National.

2.2. Varian Associates

Joseph Holloway had left National and joined the Quantum Electronics group of the Bomac division of Varian Associates in Beverly, Massachusetts. There he led development of a new, completely enclosed caesium-tube including the deflection magnets. The tube was 0.61 m long and about 0.13 m in diameter and was given the model number BLR2. The microwave cavity was constructed from an X band silver-lined E-plane waveguide bent into the shape of a flattened U with a tee-feed at the centre of the bottom of the U as shown in figure 1. The tips of the U were closed with short circuit plates to form a closed resonant cavity. Holes for the caesium-beam to pass were placed at the tips of the U, and microwave leakage from these holes was minimized by having waveguides beyond the cutoff extended for the beam to pass through. Thus, Ramsey-separated field excitation was obtained with fairly good symmetry and a consequent small differential phase-shift between the excitation regions. This led to good accuracy as will be discussed later. This tube was put in the market in the early 1960s, and a number were sold.

In about 1962, the Frequency and Time Division of Hewlett-Packard gave the Varian group a contract to develop a shortened, proprietary version of the tube, 0.406 m long. This new tube would fit horizontally, front-to-back or side-to-side in a standard electronics rack package, thus allowing substantial reduction in size of the frequency standard. This shorter tube had a larger resonance linewidth, which increased the burden on the electronics.

Hewlett-Packard purchased the Varian Bomac Quantum Electronics group and the tube manufacturing rights in 1967.

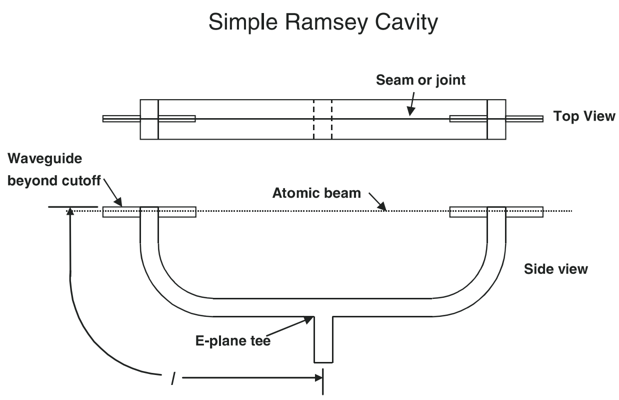

Figure 1. Schematic of a simple Ramsey cavity. Symmetry about the E-plane tee is critical. l is the length of a cavity-half. $2l$ should be an integral number of guide-wavelengths. The seam or joint is the symmetry plane for two machined cavity-halves. The microwaves are coupled into the leg of the tee with a coupling-loop or probe.

They continued to sell the longer BLR2 tube to other companies after the purchase. However, it is no longer available today.

2.3. Pickard and Burns

In the mid-1960s, Pickard and Burns developed a transistorized caesium-standard using the Varian BLR2 tube. It is not known how many of these units were sold or what their performance was. The unit was on the market for only a relatively short time.

2.4. Hewlett-Packard/Agilent Technologies

In the early 1960s, the Frequency and Time Division of Hewlett-Packard became interested in caesium-beam frequency standards. They purchased one of the 0.61 m long BLR2 tubes from Varian and developed a completely transistorized set of electronics to interface with it. Subsequently, as mentioned above, Hewlett-Packard awarded a contract to Varian to develop a 0.406 m long tube. This tube and the transistorized electronics were packaged in a standard rack height instrument 0.222 m high and put on the market in 1964 as the HP5060A caesium-beam frequency standard. It was the first caesium-beam frequency standard with all solid-state electronics [5]. It was displayed at the 1964 International Chronometry Conference in Lausanne, Switzerland and, at that conference, made one of the first high-accuracy measurements of the hydrogen hyperfine frequency, as produced by a Varian Hydrogen Maser, with respect to caesium. The HP5060A had a high-quality quartz oscillator that could maintain good performance even with the caesium-tube shut down.

The accuracy of the HP5060A originally specified in 1964 was $\pm 5 \times 10^{-11}$. It was realized that this was extremely conservative, and so, after several years experience with the standard, the accuracy and long-term stability of the product were then specified as $\pm 1 \times 10^{-11}$. The short-term stability was about $3 \times 10^{-11}\tau^{-1/2}$, where τ is the averaging time. A fair number of these standards were sold to time-keeping laboratories, other commercial applications and the military.

As mentioned above, in early 1967 Hewlett-Packard purchased the Varian Bomac Quantum Electronics group and the manufacturing rights for the caesium-beam tubes. The operation continued at its Beverly location until it was moved to Santa Clara, California in late 1969. All Hewlett-Packard (and now Agilent) caesium-tube and frequency standard manufacturing was, and still is, done at the Santa Clara location.

Improvements in the electronics, mainly as a result of advances in semiconductor technology, were incorporated in the model HP5061A, introduced in 1967. This unit also had an LED clock and a backup battery available as options. It was sold in fair quantities until 1986, when the HP5061B was introduced.

Early in 1969 work had started on an improved design of the 0.406 m tube. The caesium flux was greatly increased in order to get improved short-term stability. The magnetic shielding and field homogeneity were greatly improved by using a box type of shield with an internal baffle to help shield the large hole necessary for the microwave feed. Coils wound inside the box and around the baffle produced the small C-field required, which was very uniform over the caesium-beam path (see figure 3). A new microwave structure was designed that used, instead of a bent waveguide, a machined cavity built in halves around the symmetry plane and manufactured in such a way as to minimize asymmetries. This guaranteed that the differential phase-shift between cavity ends could be kept very small. This tube was made available as a high-performance option with improvements in accuracy and long- and short-term stability. It first appeared on the market in 1972 and was available as a high-performance option for the HP5061A.

A very small caesium-tube, model 5084, was developed in the early 1970s. It was about 0.17 m long and 0.102 m in diameter. The 5084 was packaged along with a modified set of electronics into a standard, the model 5062C, that had a small overall size, but at the expense of reduced performance. A number of these units were sold to the US Navy.

The 5061B was introduced in 1986. It had an improved backup-battery life and an LCD clock and minor changes in the electronics. It was available with the standard-performance and high-performance caesium-tubes.

In 1988, work was started on a new caesium-standard with digital electronics. This unit, model HP5071A, was microprocessor based. One of the major areas of performance that needed improvement over earlier standards was frequency change with temperature and/or humidity. The main cause of these frequency changes is a combination of cavity detuning and lack of constancy of the microwave magnetic field in the interaction regions. Variations in the small static magnetic field required, the C-field, also contributed. These effects were greatly reduced by closed-loop techniques to be described later in this paper. In addition, improvements were made on the tube and its diameter was reduced so that the instrument height was reduced from 0.222 m to the standard rack height of 0.133 m. The improved tube had a very carefully optimized set of beam optics. The HP5071A first appeared on the market in 1991 and had a performance that was a great improvement over all previous commercial caesium-standards. Present specifications for the high-performance option include an accuracy of $\pm 5 \times 10^{-13}$, a long-term stability within $\pm 1 \times 10^{-13}$, a short-term stability of $8.5 \times 10^{-12}\tau^{-1/2}$ for $100\,\mathrm{s} \leqslant \tau \leqslant 3 \times 10^{6}\,\mathrm{s}$, where τ is the averaging time in seconds, a flicker of frequency floor $<1 \times 10^{-14}$ and an almost negligible frequency variation with temperature, humidity and magnetic field.

In 1999, Agilent Technologies was split off from Hewlett-Packard as a separate company. The frequency standard business went with Agilent along with its manufacturing and research and development. The 5071A continued, and continues today, to be sold under Agilent's name.

2.5. Oscilloquartz

In 1964 the Swiss company Oscilloquartz delivered its first commercial caesium-standard to the ESRO, the predecessor of the European Space Agency (ESA). In 1988 Oscilloquartz, in a joint venture with SERCEL in France, developed a new European digital caesium-standard (EUDICS) with digital control loops and remote operation capability via RS232. The accuracy of this unit was $\pm 3 \times 10^{-12}$ for their high-performance unit. Oscilloquartz has supplied many frequency and time subsystems based on their digital caesium-standards.

2.6. Frequency Electronics, Inc.

As mentioned above, in 1969 FEI purchased the Atomic Clock Division of the National Company. In the mid-1970s FEI developed an airborne military caesium-standard for use in the low-frequency communication programme. The standards were used in the planes patrolling along the 'Dew Line'. In the late 1970s and early 1980s FEI developed a smaller and lighter-weight tube and frequency standard for the US Air Force. They also supplied some standards for the GPS satellites. In the 1980s, FEI supplied a large number of standards to the US Navy for use in nuclear submarines.

2.7. Frequency and Time Systems/Datum/Symmetricom

Bob Kern, who was previously with Hewlett-Packard, founded Frequency and Time Systems (FTS) in 1971. It was capitalized by Ebauches SA of Switzerland. The first caesium-beam tube production was in 1974, and a contract was awarded that year to design and build caesium-beam clocks for GPS satellites. In 1977 FTS put the first caesium-standard in space. In 1978, the Swiss took over FTS completely and, as a result, work for the US government stopped. The year 1979 saw the first caesium-standards both as self-contained modules and as rack mounted and portable instruments.

Datum acquired FTS in 1983, with FTS retaining its name. That same year, FTS brought out the industry's first microprocessor controlled caesium-standard. Production contracts were awarded for GPS Block II and IIA Satellite caesium-standards from 1983 through 1986.

Symmetricom acquired FTS in 2002, and the name was changed to Technology Realization Centre.

The company has a number of caesium-standards on the market including several high-performance units (4065C, CsIII-EP and Cs4000-EP) that have performance specifications essentially the same as those of the Agilent 5071A.

2.8. Kernco

Bob Kern, who had left FTS, incorporated Kernco in 1978. The company was devoted to space work. In 1979 Kernco was awarded a contract by the Naval Research Laboratory to develop a GPS flight caesium-standard. From 1980 to 1988 Kernco built several second-source flight units. Units were flown on satellites 29, 30 and 34 and, as reported by the Naval Research Laboratory on 29 October 2002, were the best performing caesium-standards in the GPS satellite constellation.

2.9. Others

There has been work in Japan on laboratory caesium clocks. It is believed that there was some work on commercial clocks, but the extent of that work and whether there have been sales is not known.

A number of years ago, there was a Russian commercial caesium clock put on the market by the 'Quartz' Research and Production Association, Gorkii, USSR. This is the same group that produced hydrogen masers. This caesium-standard did not stay on the market for long and it is not known what the specifications were.

3. Technical aspects

As mentioned above, important performance characteristics of commercial caesium-standards are accuracy, short- and long-term stability and frequency dependence on the environment. Since the commercial tubes are short compared with laboratory tubes, the resonance linewidth is greater. Consequently, great care must be taken with regard to line asymmetries and electronics-induced resonance offsets. A number of the most important technical details and design features that address these performance issues are given in this section of the paper.

3.1. Energy levels

The energy levels involved in caesium-beam frequency standards are those present in the ground state due to the hyperfine-interaction splitting. Because the nuclear spin of ^{133}Cs is $\frac{7}{2}$ and there is a single electron in an S state outside closed shells, the total angular momentum in the hyperfine states is $F = 3$ or 4. These two angular momentum states are separated in energy at zero magnetic field by $\delta E = h\nu_{00}$, where h is Planck's constant and ν_{00} is the hyperfine resonance frequency, 9192 631 770 Hz. In the presence of a small magnetic field these levels are split into $2F + 1$ levels, seven for $F = 3$ and nine for $F = 4$. The levels used for frequency standards are $F = 3$, $M_F = 0$ and $F = 4$, $M_F = 0$, where M_F is the quantum number of the projection of the angular momentum onto the direction of the magnetic field. These levels are used because, for small magnetic fields, their energy depends only on the square of the field strength, whereas all the other states have first order dependence. The hyperfine resonance frequency for the transitions between these frequency standard levels in the presence of an applied small static magnetic field B is designated ν_0, where $\nu_0 = \nu_{00} + KB^2$, with $K = 427.447 \times 10^8\,\mathrm{Hz\,T^{-2}}$. When a microwave magnetic field at the resonance frequency, ν_0, is applied, transitions between these levels are induced. This transition will be referred to as the 0–0 transition. The applied static B field is designated the C-field, and its function is to separate in frequency all the possible transitions by a small amount so that a clean 0–0 transition can be observed. The fractional frequency offset due to a C-field of 6 μT is $(\nu_0 - \nu_{00})/\nu_{00} = 1.674 \times 10^{-10}$ and is taken into account for the actual output frequency of the standard.

At room temperature, $h\nu_0 \ll kT$, where k is Boltzmann's constant and T is the absolute temperature. As a result, all the ground state hyperfine levels are almost equally populated and so there are about an equal number of induced transitions between the lower and upper groups of states in either direction. Some form of state selection must be used to unbalance the populations and detect the transitions. In all the present commercial caesium-beam tubes, strong inhomogeneous magnetic fields are used to deflect the atomic beam for state selection and detection. The $F = 4$ states with the exception of $M_F = -4$ are deflected away from strong fields while the $F = 3$ states and the $F = 4$, $M_F = -4$ state are deflected towards strong fields. A simplified schematic of a magnetically deflected beam-tube is shown in figure 2.

3.2. Transition probability

The steady-state transition probability for Ramsey excitation for a single atomic velocity and uniform magnetic C-field is

$$P = \frac{b^2 \sin(2pt)^2}{p^2}\left(\cos\left(\Delta T + \frac{\theta}{2}\right) - \frac{\Delta}{p}\sin\left(\Delta T + \frac{\theta}{2}\right)\tan(pt)\right)^2, \qquad (1)$$

where $\Delta = (\omega - \omega_0)/2$, $\omega_0 = 2\pi\nu_0$, ω is the angular microwave excitation frequency, t is the transit time through a cavity end, T is the transit time in the drift region between the cavity ends, θ is the phase angle lead of the microwaves in the second cavity end with respect to that in the first, $2b$ is the Rabi frequency of the caesium in the microwave magnetic

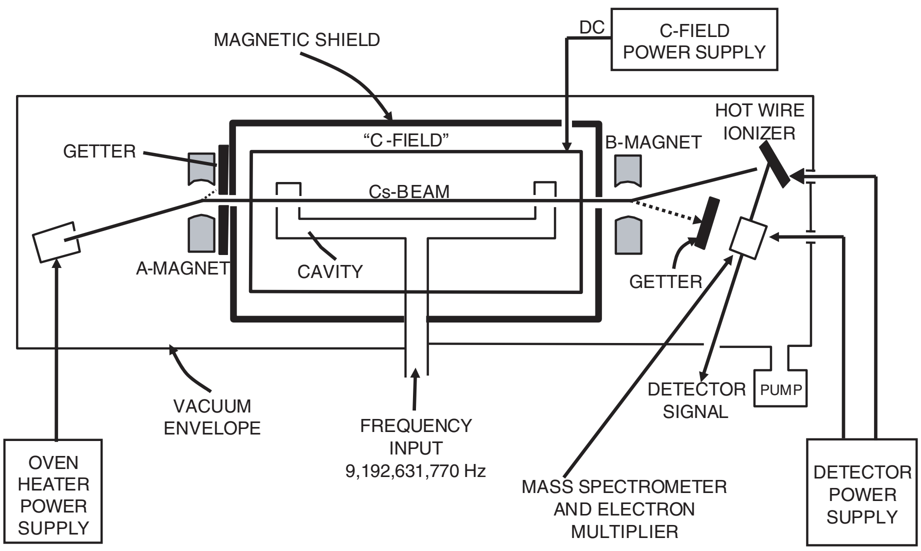

Figure 2. Schematic of a magnetically deflected caesium-tube. The atoms that enter the cavity from the oven are in the $F = 3$ state since that state is deflected towards stronger fields. The deflection angles are greatly exaggerated. The C-field and the microwave magnetic field are perpendicular to the page. The coil for generating the C-field would be wound inside the magnetic shield with its axis perpendicular to the page. The dimension d (see figure 3) is also perpendicular to the page.

field and $p = (\Delta^2 + b^2)^{1/2}$. Note that b is proportional to the microwave magnetic field amplitude in the cavity ends. It is assumed that the microwave amplitude in the cavity ends is uniform, so that the atoms see square pulses of microwaves. The rotating field approximation has been used and other, higher energy, levels have been neglected. Otherwise, the result is an exact solution and is a very good approximation. Equation (1) is equivalent to Ramsey's equation (V.37) with different notation [6]. The actual output current produced by the beam-tube is proportional to the transition probability averaged over the velocity distribution in the detected beam plus some background. That ω for which the probability is maximum is one measure of the output frequency of the standard. For resonance lines that have even symmetry about $\Delta = 0$, such as equation (1) when $\theta = 0$, perfect modulation schemes to find the line centre will give the same result as the position of the line maximum.

For $\theta = 0$, the smallest value of b for unity transition probability when $\Delta = 0$ in the single velocity case is $b = \pi/4t$. Then, for $\Delta \ll b$, a good approximation for the transition probability is

$$P \cong \cos\left(\Delta T\left(1 + \frac{4t}{\pi T}\right) + \frac{\theta}{2}\right)^2. \quad (2)$$

From this, the full linewidth in hertz at half maximum for $\theta = 0$ is

$$\text{linewidth} \cong \frac{1}{2T\,(1 + 4t/\pi T)}. \quad (3)$$

If the average of B^2 in the cavity ends differs from the average of B^2 in the drift region, the corresponding resonance frequencies are not equal and the steady-state transition probability becomes

$$P_1 = \frac{b^2 \sin(2p_1 t)^2}{p_1^2}\left(\cos\left(\Delta T + \frac{\theta}{2}\right) - \frac{\Delta_1}{p_1}\sin\left(\Delta T + \frac{\theta}{2}\right)\tan(p_1 t)\right)^2, \quad (4)$$

where $\Delta_1 = (\omega - \omega_1)/2$, ω_1 is the angular hyperfine resonance in the cavity ends, $p_1 = (\Delta_1^2 + b^2)^{1/2}$ and $\Delta = (\omega - \omega_0)/2$, with ω_0 being the average angular hyperfine resonance in the drift region between the cavity ends. The author derived equation (4) using a method similar to Ramsey's derivation of his equation (V.37) [6].

The transition probability for unequal resonance frequencies at the cavity ends as well as the drift region has also been calculated, but the results are quite complicated, unpublished and will not be reproduced here.

3.3. Error due to differential phase, θ

Equation (1) shows that the maximum probability occurs when $\Delta \cong -\theta/2T$ for $t \ll T$. For a typical Agilent high-performance commercial beam-tube, the operating parameters are approximately $t = 80 \times 10^{-6}$ s, $T = 1.3 \times 10^{-3}$ s and $b = 9817\,\text{s}^{-1}$. For the phase-shift induced magnitude of the fractional frequency error to be less than 1×10^{-13}, $|\theta|$ must be less than 7.5 µrad. A calculation for the differential phase-shift at the ends of a low loss waveguide cavity due to unequal lengths of the cavity-halves gives [7]

$$|\theta| \cong \alpha l k_0 |\delta l|, \quad (5)$$

where α is the waveguide attenuation constant, l is the length of a cavity-half (see figure 1), which is ideally an integral number of half guide-wavelengths, $k_0 = 2\pi/\lambda_g$, with λ_g the guide-wavelength, and $|\delta l|$ is the difference in effective length of the two cavity-halves. For a low-loss waveguide cavity typical of that in a high-quality commercial beam-tube, $|\delta l| = 25$ µm gives about 7 µrad for the phase difference. It is assumed that the attenuation is uniform throughout the cavity and that the waveguide dimensions are constant throughout the cavity structure. If these are not satisfied, the phase-shift could be larger. These considerations indicate that very careful manufacturing is required for the cavity.

Since the resonance line is asymmetric when there is differential phase-shift, it may be possible to estimate the phase-shift by careful measurements of the asymmetry. Knowledge of the velocity distribution of the detected atoms is also necessary for this estimation.

The distributed phase-shift at the cavity ends is very difficult to calculate for the waveguide cavity structures. Some comments can be made. If the waveguide ends have even mechanical symmetry about the centre of each end along the beam path, the distributed phase-shift must also have even symmetry. Atoms encountering such a symmetric distributed phase-shift will see instantaneous changes in frequency on going through a cavity end but no shift in the average frequency. However, there is an angular spread to the beam, and so the atoms may experience some asymmetry. It is not known how big this effect is.

3.4. Error due to inhomogeneity in the C-field

The value of ν_0 is determined mainly by the average value of B^2 along the beam path in the drift region between the cavity ends. What can be measured, however, is the average of B. Measuring the Zeeman frequency does this. $\overline{B^2}$ is equal to $\bar{B}^2$ (where the bars indicate the average of the quantity) only if the field is homogeneous. In fact, the variance of the field is $\sigma^2 = \overline{B^2} - \bar{B}^2$. Since $\sigma^2 \geqslant 0$, $\overline{B^2} \geqslant \bar{B}^2$. Thus an inhomogeneous C-field leads to an underestimate of the shift of ν_0 from ν_{00}. A $\pm 1\%$ error in the value of $\overline{B^2}$ gives a frequency error of $\pm 3.36 \times 10^{-12}$ at a C-field of 6 µT. The design and construction of the magnetic shields and C-field forming the structure are thus very important.

3.5. Error due to difference between magnetic fields of the cavity ends and drift region

Equation (4) gives the transition probability for the case where $\overline{B^2}$ differs in the drift region and the cavity ends, causing a difference in the hyperfine resonance frequencies in these regions. The exact expression for the frequency at which the maximum occurs is complex. A numerical solution for the case where the rms fields in the regions differ by 1% in a typical commercial tube gives a fractional frequency offset of about 2×10^{-13} from the hyperfine resonance frequency in the drift region. The offset is linear in the percentage rms field difference between the regions. This effect points out again that careful shield design and construction are needed to prevent leakage from the strong deflecting-fields affecting the C-field at the cavity ends.

3.6. Frequency offset due to adjacent field dependent transitions (Rabi and Ramsey pulling)

As mentioned above, there are other microwave transitions nearby in frequency when a static magnetic field, the C-field, is present. These are the Zeeman transitions. If the microwave magnetic field and the C-field are perfectly parallel, only those transitions between the states with identical M_F are allowed. For example, transitions between $F = 3$, $M_F = +1$ and $F = 4$, $M_F = +1$ are allowed. There are seven such possible transitions, and if the microwave frequency is scanned, all of them can be observed. Their separation in frequency is approximately $7\,\text{kHz}\,\mu\text{T}^{-1} \times B$, where B is the average C-field in the drift region. This is about 42 kHz in a 6 µT C-field. Because t and T are finite, the transitions have finite width and they overlap the 0–0 frequency-standard transition to a small extent. If the velocity distribution is wide enough, the oscillatory behaviour of these adjacent transitions can be fairly well damped by the velocity averaging, so that what remains at the 0–0 transition frequency is a small transition probability slope that can shift the maximum of the 0–0 frequency standard transition. Only the closest transitions need be considered. If the transition probabilities for the above and below transitions are the same, the slopes will cancel. Unfortunately, for magnetic deflection, this is not the case. The fractional frequency-shift of the 0–0 transition, called Rabi pulling, is then given by

$$\left|\frac{\delta\nu_0}{\nu_0}\right| \cong \frac{1}{64\pi^2T^2t^2\nu_0\nu_z^3}\left|\frac{\delta P}{P}\right|, \qquad (6)$$

where ν_z is the Zeeman frequency (about 42 kHz), and $\delta P/P$ is the fractional difference in peak probability between the upper and lower Zeeman transitions. For a 10% imbalance in the probabilities and a conventional beam-tube with $\nu_z = 42$ kHz, $t = 80 \times 10^{-6}$ s, $T = 0.0013$ s and $\nu_0 = 9192$ MHz, $|\delta\nu_0/\nu_0| \cong 2.1 \times 10^{-14}$.

The 0–0 transition is pulled towards the larger of the adjacent transition probabilities.

If the velocity averaging is insufficient to reduce the oscillatory behaviour to a negligible amount, the effects can be much larger.

The frequency-shift caused by the linear background slope can be removed by using sinusoidal modulation with third harmonic detection. It can also be removed by using slow interrogation at four frequency points: $\nu 1 = \nu + \nu_i$, $\nu 2 = \nu - \nu_i$, $\nu 3 = \nu + 3\nu_i$ and $\nu 4 = \nu - 3\nu_i$. ν_i is chosen to be about the magnitude of the frequency difference between the first inflection point and the maximum of the transition probability. It is easy to show that if the error signal is then taken to be

$$\text{Errsig} = 3[I(\nu 1) - I(\nu 2)] + [I(\nu 4) - I(\nu 3)], \qquad (7)$$

where $I(\nu 1)$ is the beam-tube output current at frequency $\nu 1$, etc, then any linear background slope does not affect the error signal. The value of ν that makes the error signal vanish is exactly that at which the probability maximum occurs if the line, without background slope, has even symmetry.

When the microwave magnetic field is not completely parallel to the static C-field, transitions where M_F changes by ± 1 are allowed. These Zeeman transitions are midway in frequency between those for which M_F does not change. The transitions that are closest to the 0–0 transition involve transitions to or from one or the other of the $M_F = 0$ clock states. This requires that these pulling effects on the 0–0 transition, called Ramsey pulling, be treated quantum mechanically instead of just considering a slope [8]. The analytical results of the calculations are quite complex and will not be presented here. Calculation shows that for an angle of 5° between the C-field and the microwave magnetic field the magnitude of the pulling can be as large as 2×10^{-13}. The pulling is proportional to the fractional imbalance in the adjacent transitions where M_F changes by ± 1. Ramsey pulling differs from Rabi pulling in that it cannot be removed by using sinusoidal modulation and third harmonic detection or the four-point interrogation technique mentioned above. Careful design and construction of the microwave cavity and the C-field structure and shields are required.

3.7. Frequency offset due to spectral impurities in the microwave interrogation signal

The microwave signal that interrogates the atoms is usually derived by frequency multiplication and synthesis from a quartz crystal oscillator. As a result, the microwave signal contains other frequencies besides the main interrogation signal. These can cause frequency offsets and the effect has been well treated in the literature [9, 10]. An expression for the fractional frequency offset of the maximum transition probability caused by a small signal outside the range where all the Zeeman transitions occur is

$$\frac{\delta\nu}{\nu_0} = \frac{4b_1^2}{\omega_0(\omega_0 - \omega_1)}\frac{l}{L}\frac{\tan(bt)}{bt}, \qquad (8)$$

where $2b$ and $2b_1$ are the Rabi frequencies of the main and small signals, respectively, and ω_0 and ω_1 are the unperturbed hyperfine angular resonance frequency and the angular frequency of the small perturbing signal, respectively. t is the transit time through the cavity end microwave interaction region, and l and L are the lengths of the cavity end microwave interaction regions and the drift length, respectively. $bt \cong \pi/4$. Putting in the values for a typical high-performance tube, the offset is about 2×10^{-13} for a perturbing signal 150 kHz away and down 40 dB from the main signal. Velocity averaging will affect this to some degree. The offset is linear with the power ratio of the perturbing signal to the main signal provided the main signal is optimized to give approximately maximum transition probability. Equation (8) shows that the pulling varies inversely with the frequency difference from the main signal. Note that if there are two perturbing signals, outside the Zeeman range, equally spaced above and below the main signal, their pulling effects are opposite in sign and the total pulling will be zero if their amplitudes are equal.

For perturbing frequencies within the Zeeman region but still outside the 0–0 Rabi pedestal, the pulling will be less than 1.5×10^{-13} if the perturbing signal is 50 dB below the main signal for a typical high-performance tube. This result includes the case where the perturbing signal frequency is on a point of maximum slope of one of the Zeeman transitions. This case is a classical effect where the slope of the Zeeman transition must be cancelled by the slope of the main transition.

When the perturbing signal lies within the main Rabi pedestal, the situation is much more complicated. The analytic results are given in [10]. The pulling oscillates with the same period as the Ramsey pattern. For a tube similar to a typical high-performance tube, the worst case pulling for a perturbing signal down 70 dB from the main signal is about 1.5×10^{-13}.

In general, two equal-amplitude small perturbing signals spaced exactly the same frequency above and below the main signal frequency cause zero pulling in all cases except when they lie directly on a pair of the interrogation frequency modulation sidebands. It is easy to show that this situation is equivalent, in the general case of an arbitrary initial phase between the perturbing signals, to a combined amplitude and phase modulation of the main signal at the difference frequency between the frequency of one of the perturbing signals and the main signal frequency. These modulations are either in phase or π radians out of phase. The phase modulation is equivalent to frequency modulation shifted by $\pi/2$ radians, so the amplitude and frequency modulations are in quadrature. Pure amplitude or frequency modulations are special cases. When the frequency difference between one of the perturbing signals and the main signal is much less than the Ramsey linewidth, the effective modulation can be considered slow. All these types of slow modulated signals applied to a symmetric resonance line will cause no offset in the time-averaged frequency. Thus, in all cases of two equal-amplitude perturbing signals, with the exception of when they lie directly on intentional modulation sidebands, there is no frequency offset.

In contrast, a single, small perturbing signal is equivalent to a combination of amplitude and frequency modulations of the main signal that are in phase. Consider a very small frequency difference between the perturbing and main signals so that the modulation can be considered slow. This always causes an offset of the frequency of the beam-tube current maximum when applied to a symmetric resonance line except when there is a nonlinear dependence of the transition probability on the microwave amplitude such that it has a maximum. This exception is the case with a caesium beam-tube as can be seen from equation (1), where the transition probability for a single beam velocity has a maximum for $bt = \pi/4$ when $\Delta = 0$. This is also true when there is a beam velocity distribution with perhaps a different value of b. If the average microwave amplitude is at the value that gives maximum probability, slow changes in amplitude will result in only second order changes in beam-tube current so that the effective amplitude modulation is greatly reduced and is at the second and higher even harmonics. Thus, the offset caused by a close, small perturbing signal can be greatly reduced by this technique. A servo loop that keeps the probability at its maximum by adjusting the microwave amplitude is easy to implement (see section 3.13).

3.8. Frequency offset due to even order distortion in the interrogation modulation

Shirley has done a thorough calculation on the frequency offset due to distortion in slow sinusoidal interrogation modulation [11]. One result is that even-order distortion generally causes frequency offset. Calculation for a typical high-performance tube shows that if the second harmonic is about 90 dB below the fundamental at the optimum modulation index and synchronous detector reference phase, the worst case fractional frequency offset will be about 1×10^{-13}. The value of 90 dB is fairly easy to achieve by starting with a frequency twice the desired modulation frequency signal and using it to trigger a flip-flop. If this is done carefully, the flip-flop output will be a square-wave containing only a very small amount of even harmonics. A linear, passive filter removes most of the remaining odd harmonics, leaving a fairly clean sine wave that can be applied to a linear, low-level phase modulator early in the frequency multiplier chain. This type of sine wave generation for interrogation modulation was used in many early caesium-beam standards.

Another important result is that odd-order distortion, such as exists in square-wave frequency modulation, causes no offset. It is easy to show that the best slow frequency modulation for interrogation is square-wave frequency modulation since it can give the largest error signal for a given average microwave input frequency offset. An ideal square-wave contains only odd harmonics. It is easy to generate square-wave frequency modulation by switching a synthesizer that is part of the microwave frequency chain. See section 3.13 for further discussion.

3.9. Microwave cavity considerations

From considerations of differential phase-shift discussed in section 3.3, it is desirable to have α in equation (5) as low as possible. This implies a high Q_{cavity}, where Q is the usual quality factor for a resonant structure. However, cavity detuning from ν_0, caused by a temperature change, for example, can produce a frequency offset because the applied slow-frequency-modulated signal then has in-phase amplitude modulation due to the slope of the detuned cavity response. Consider the case of slow square-wave frequency modulation. Then the effect causes a frequency offset with magnitude proportional to $(Q_{\text{cavity}}/Q_{\text{atomic line}})^2 \times |\delta A \times \delta\omega_c|$, where δA is the departure in microwave amplitude from the point that gives maximum beam-tube output current at the frequency modulation peaks and $\delta\omega_c$ is the cavity detuning. This is similar to the single small perturbing signal case in section 3.7, but the frequency modulation here is intentional. The cavity detuning effect thus calls for a low Q_{cavity}. Both these desires for Q_{cavity} can be met by having the cavity made from low-loss material, perhaps silver plated, and introducing a loss only in the waveguide tee-feed. Then the low-loss cavity-halves have a high effective Q_{cavity}, while the overall Q_{cavity} can be made very low. Figure 1 is a schematic diagram of a microwave cavity typical of those used in commercial beam-tubes.

The frequency-shift caused by the cavity detuning can be greatly reduced by servo control of the microwave amplitude to be at the value that gives maximum transition probability. It can be shown that this technique works best, by far, when slow-square-wave frequency modulation is used. This is discussed further in section 3.13.

Maintaining high mechanical accuracy and the use of good, low-loss materials is important from the standpoint of differential phase-shift as mentioned earlier.

The caesium-beam should go through the cavity ends as close as possible to and the same distance from the end shorts to minimize differential phase-shifts. Also, the beam should go through the centre of the long dimension of the waveguide cross section to have the microwave and static fields as parallel as possible to minimize Ramsey pulling.

It is very important that the whole microwave structure including the cavity, feed cables and the microwave-generating source be well shielded and free from microwave leaks so that the atoms see the microwaves only at the cavity ends. The presence of fields elsewhere, particularly in the drift region, can cause large frequency offsets.

3.10. C-field and magnetic shield structure

The functions of this structure are to shield the interaction and drift regions from external magnetic fields and to provide the small, uniform C-field in these regions. Different manufacturers use different structures and different numbers of shields. Figure 3 shows one construction technique used for obtaining these functions. Shown is a box structure surrounded by two outer shields, all of high permeability. The coil is wound inside the box as close as possible to the sidewalls. The distances of the top and bottom coil windings from the top and bottom of the box are half the pitch of the rest of the coil. If the box has infinite permeability, the coil is reflected in the top and bottom and so the region inside effectively sees an infinitely long, uniform solenoid, which has an almost uniform field everywhere inside except very close to the coil. For this to be true, the distance, d, between the top and bottom of the box must be constant. The local field variations are given by $\delta B/B \cong -\delta d/d$ for regions about the size of d.

The outermost shield must be thick enough not to saturate with the highest external field expected. The box walls must be thick enough not to saturate when carrying the flux necessary for the C-field.

3.11. Relativistic effects

The lowest-order relativistic effect for a stationary caesium-standard at sea level is due to the caesium-beam velocity. The fractional frequency-shift is $\delta\omega/\omega_0 \cong -V^2/2c^2$, where c is the speed of light and V is the rms speed of the caesium-beam. For a typical high-performance tube the shift is about -1×10^{-13}. This effect is large enough for it to be taken into account in the frequency output of the standard. The distribution of beam speeds leads to a small line asymmetry.

In addition to the velocity effect, there is an effect due to the gravitational potential at the location of the standard. For a stationary standard operating at a height h above sea level, small compared with the earth's radius, the frequency is higher than that of one at sea level by a fractional amount $\delta\omega/\omega_0 \cong gh/c^2$, where g is the local measured value of the gravitational acceleration. This amounts to about $+1.09 \times 10^{-13}$ for $h = 1$ km.

Clocks in motion relative to the earth require special treatment because the earth is not an inertial frame since it is rotating. A good approximation giving the fractional frequency difference between two moving standards is

$$\frac{\nu_1 - \nu_2}{\nu_2} \cong \frac{g_1 h_1 - g_2 h_2}{c^2} - \frac{2\omega R(V_{1e}\cos\theta_1 - V_{2e}\cos\theta_2) + V_1^2 - V_2^2}{2c^2}, \tag{9}$$

where g_1 and g_2 are the measured local accelerations of gravity for clocks 1 and 2, h_1 and h_2 are their heights above sea level, V_1 and V_2 are their speeds with respect to the earth's surface, V_{1e} and V_{2e} are the eastward components of their velocities, θ_1 and θ_2 are their latitude angles, ω is the angular rotation rate of the earth, R is the radius of the earth and c is the speed of light. Equation (9) is easily derived by making the calculations from a non-rotating system and using the non-relativistic addition of velocities.

Assuming clock 2 is stationary at sea level and clock 1 is flying at an altitude of 10 km at 800 km h^{-1} at the equator in the eastward direction, the fractional frequency difference is -3.3×10^{-13}. If clock 1 is going westwards, the difference is $+19.6 \times 10^{-13}$. Assume clock 1 is set to clock 2's time at the beginning of a flight around the world. Comparing time

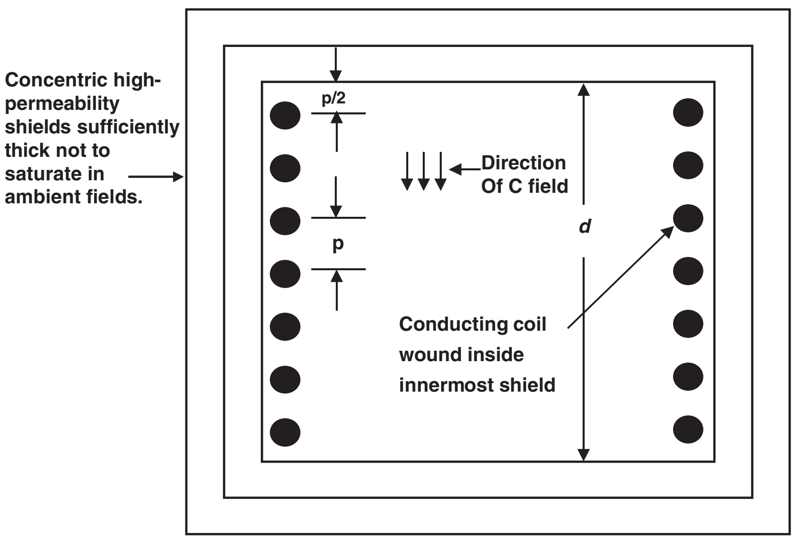

Figure 3. Sketch of one type of C-field and shield structure. With the coil spacings, p, shown, the magnetic field inside the C-field region will be very uniform if d is uniform.

after the flight, clock 1 will be about 60 ns behind clock 2 for an eastward flight and about 350 ns ahead for a westward flight. The different results for east versus west are due to the earth's rotation. Hafele and Keating [12] did such a flying clock experiment in the early 1970s. Taking clock 1 and transporting it around the earth's equator at an infinitesimally small speed at sea level gives ± 206.5 ns for the time difference, positive for westwards.

Since commercial caesium clocks are often transported, these effects must be taken into account for the highest accuracy results.

3.12. Short-term stability

The short-term stability in well-designed beam-tubes and electronics is limited by the shot-noise in the detected caesium-beam current. Ideally, the electron multiplier and following amplifiers should add only small amounts of noise. Present high-performance standards have a short-term stability of about $8.5 \times 10^{-12}\tau^{-1/2}$, with τ the averaging time in seconds over the appropriate range of time.

3.13. Electronics

Typically, a high-quality 5 MHz or 10 MHz quartz oscillator with electronic frequency control is frequency multiplied by an integer N to a frequency close to the caesium resonance and the difference is made up with a high-resolution synthesizer to produce the microwave interrogation signal frequency at ν_0. The synthesizer produces a chosen rational fraction, F, of the oscillator frequency. A low frequency error signal is derived through frequency modulation of the microwave signal and synchronous detection of the beam-tube output. The modulation is often obtained by switching the frequency of the synthesizer. Double integration (paralleled with single integration for loop stability) of the error signal and feeding the result to the electronic frequency control input of the oscillator closes the servo loop and the oscillator frequency is locked exactly to the value $\nu_{osc} = \nu_0/(N + F)$, provided the frequency multiplier and synthesizer are phase stable with time, temperature, etc. The purpose of the double integration is to completely remove the linear frequency drift of the oscillator. Also, with double integration, if the oscillator has a step-change in its free-running frequency, the integral over time of the change in the frequency standard's output frequency is zero, resulting in no long-term time or phase error caused by the step.

Since the frequency that is really tightly controlled is that of the microwave input to the beam-tube, if the multiplier and/or synthesizer have dynamic phase-shifts, the oscillator phase, and consequently frequency, will change dynamically to counteract the shifts. Rather than use the oscillator signal as the standard's output, it is better to take a signal out of the frequency multiplier at some high multiple and frequency divide it to get the output signal. Frequency dividers can be made with better phase stability than multipliers.

As mentioned earlier, the best signal to noise ratio is obtained by using slow square-wave frequency modulation of the microwave signal. The peak frequency deviation should reach the first inflection points of the transition probability. The modulation can be done accurately and easily by programming a digital frequency synthesizer used for the required synthesis at the top of the frequency multiplier chain. The synchronous detection of the beam-tube signal and all the integration are best done digitally with the digital output fed to a digital-to-analogue converter. The synchronous detection can be done by using a voltage-to-frequency converter on the amplified beam-tube signal and counting its output frequency after the transient from the frequency change has died out and before the next transient starts. The differences between the counts from the two sides of the transition probability form the digital error signal. Great care must be taken to avoid any leakage of the modulating square-wave into the circuitry ahead of the voltage-to-frequency converter or its power supply. Such leakage results in a false error signal that can cause a large frequency offset.

As mentioned in section 3.9, cavity detuning coupled with departure of the microwave amplitude from its value that maximizes the transition probability leads to a frequency offset when slow-modulation-interrogation is used. When the cavity is detuned, it has a slope in its frequency response. The slow frequency modulation then produces amplitude modulation in the cavity. This modulation is in phase with the desired frequency modulation and can produce fairly large frequency offsets. Consider square-wave frequency modulation. A good way to minimize the frequency-shift is to control the microwave amplitude to be at the maximum of the transition probability as a function of the microwave amplitude at the peaks of the frequency modulation. This can be done easily by slowly dithering the amplitude and synchronously detecting the resulting beam-tube output current to develop an error signal that feeds the amplitude controller to drive the amplitude to the peak, where the error signal vanishes. Doing this greatly reduces all effects due to slow, in-phase amplitude and frequency modulation. As a result, frequency changes due to temperature variation, etc, are reduced to insignificant values. This technique is the main reason why present commercial caesium-beam standards have such good performance with respect to temperature and humidity variations.

The C-field also can be servo controlled by measuring the frequency of one of the Zeeman transitions and setting it to the desired value by controlling the C-field current. This only sets the average of B while what is desired is the average of B^2, and so, as mentioned earlier, good uniformity of the C-field is necessary.

Many of the functions mentioned here plus others are implemented by a microprocessor in most of today's commercial caesium-standards.

4. Possible future directions for commercial caesium-standards

The major shortcoming in present caesium-standards is their short-term stability. Magnetically deflected beam-tubes have just about reached their limit in detected beam intensity, which determines the shot noise limitation. Optically pumping a caesium-beam with laser light for state selection and detection can improve the situation. Much more of the beam emanating from the caesium oven can be used since the velocity dependence of the magnetic state-selection is gone.

In addition, properly pumping with two laser frequencies can put essentially all the atoms in one of the clock states instead of $\frac{1}{16}$ of them, the case in magnetic state-selection. These features will improve the short-term stability. Getting a low-noise, narrow-linewidth, long-lifetime solid-state laser at the right wavelength for the small quantities required is not easy, but not impossible. Finally, there should be improvements in accuracy and long-term stability due to the improved C-field homogeneity resulting from the absence of the close-by, high-field deflecting magnets.

Another possible future direction is to make a commercial version of the caesium fountain [13]. The linewidth is proportional to (toss height)$^{-1/2}$ and so not a lot is lost by going to a fairly small apparatus. However, the laser setup necessary for a fountain is complex. The sensitivity of a fountain to acceleration is high and so its use might be restricted to the laboratory.

Acknowledgments

The author thanks the following people for information and helpful discussions: Robert Kern of Kernco Inc., Mike Garvey of Symmetricom, Marvin Meirs of FEI and Miao Zhu, Jim Johnson and Curt Flory of Agilent Labs. The author also thanks the referees for a number of corrections and suggestions that improve the accuracy and readability of this paper.

References

[1] Forman P 1985 *Proc. IEEE* **73** 1181–204
[2] McCoubrey A 1996 *Proc. Int. Frequency Control Symp. (Honolulu, HI)* pp 1225–41
[3] Beehler R 1967 *Proc. IEEE* **55** 792–805
[4] Ramsey N 1963 *Molecular Beams* (Oxford: Clarendon) p 124
[5] Bagley A and Cutler L 1964 *Proc. Int. Frequency Control Symp. (Atlantic City, NJ)* pp 344–65
[6] Ramsey N 1963 *Molecular Beams* (Oxford: Clarendon) p 128
[7] Lacey R 1968 *Proc. Int. Frequency Control Symp. (Atlantic City, NJ)* pp 545–58
[8] Cutler L, Flory C and Giffard R 1991 *J. Appl. Phys.* **69** 2780–92
[9] Shirley J 1963 *J. Appl. Phys.* **34** 783–8
[10] Audoin C, Jardino M, Cutler L and Lacey R 1978 *IEEE Trans. Instrum. Meas.* **27** 325–9
[11] Shirley J 1963 *J. Appl. Phys.* **34** 789–91
[12] Hafele J and Keating R 1972 *Science* **177** 166–70
[13] Jefferts S, Meekhof D, Hollberg L, Lee D, Drullinger R, Walls F, Nelson C, Levi F and Parker T 1998 *Proc. Int. Frequency Control Symp. (Pasadena, CA)* pp 2–5