

Accuracy evaluation of the primary frequency standard NIST-7

J. H. Shirley, W. D. Lee and R. E. Drullinger

Abstract. We develop a comprehensive approach to the evaluation of the uncertainty of a primary frequency standard. Indirect, model-dependent means are employed for the evaluation of all known biases. This approach leads to a lower combined standard uncertainty (CSU) and leaves frequency measurements under altered conditions as a powerful search technique for unknown biases. We report the results of this evaluation technique applied to NIST-7, one of the United States' primary frequency standards. The best CSU value we have reported to the Bureau International des Poids et Mesures (BIPM) is 5×10^{-15} , while a rigorous application of this evaluation process would suggest that a limiting Type B uncertainty of 3.5×10^{-15} is possible. We discuss this difference and the implications of reporting optimistic values. We also discuss the history of reported evaluations of NIST-7 including those made before the full implementation of the techniques presented here.

1. Introduction

The SI second is defined to be “the duration of 9192631770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium 133 atom” [1]. A *primary frequency standard* is a device that realizes this definition. The realization is limited by noise in the frequency measurements and by biases causing the realization to deviate from the ideals implicit in the definition [2]. Biases arise from the environment in which the atomic transition is measured, e.g. electric and magnetic fields, and motional and gravitational effects, as well as from imperfections in the apparatus.

Over the years accuracies of frequency standards have been evaluated with increasing sophistication. The capabilities of our previous standards NBS-4 and NBS-6 were described over twenty years ago [3, 4]. Preliminary evaluations of NIST-7, our first optically pumped frequency standard, have also been published [5, 6]. Uncertainty evaluations of optically pumped standards in other laboratories include those from the Laboratoire Primaire du Temps et des Fréquences (LPTF, France) [7], the National Research Laboratory of Metrology (NRLM, Japan, now the National Metrology Institute of Japan, NMIJ) [8], and the Communications Research Laboratory (CRL, Japan) [9]. Notable evaluations of magnetically state-selected standards have been reported by the Physikalisch-Technische Bundesanstalt (PTB, Germany) [10, 11].

The goal in the design and operation of NIST-7 was to realize the SI second with an uncertainty of less than 10^{-14} . The atomic line Q achieved with NIST-7 is about 1.2×10^8 . To meet our target uncertainty requires that we find the centre of the resonance line with a fractional uncertainty of less than 6×10^{-7} . This unprecedented splitting of the linewidth has motivated us to proceed with techniques and attention to details not previously practised. Although the basic hardware of NIST-7 has remained the same through the years, we have made changes in the electronics and our evaluative techniques that reduce our uncertainties. This paper describes the evaluative process we now use to measure the atomic resonance and to determine its biases and residual uncertainty.

1.1 Uncertainty evaluation

1.1.1 What we measure

We choose to operate NIST-7 intermittently as a frequency standard, rather than continuously as a clock. Periods of frequency measurement are interspersed with periods of bias determination. Frequency measurements are made by comparison with a high-quality reference oscillator, typically a hydrogen maser. We represent the result of such a comparison by [6]

$$\Delta\nu = \nu_{\text{hfs}} + \sum_j \delta\nu_j - \nu_{\text{ref}} \pm \sigma. \quad (1)$$

In this representation $\Delta\nu$ is the measured frequency difference, ν_{hfs} is the unperturbed caesium hyperfine separation defining the second, $\sum_j \delta\nu_j$ is the sum

J. H. Shirley, W. D. Lee and R. E. Drullinger; National Institute of Standards and Technology (NIST), Time and Frequency Division, Boulder, Colorado, USA.

over frequency biases, and ν_{ref} is the frequency of the reference oscillator (multiplied to the caesium frequency) during the measurement period. The uncertainty σ quantifies the statistical noise of the measurement process. We distinguish biases from noise by requiring what we call noise to have zero mean. Any effect with non-zero mean belongs in the summation of biases. In practice, the separation of σ from $\sum_j \delta\nu_j$ is never perfect. The determination of the $\delta\nu_j$ always includes some noise, while small unknown biases may remain hidden in σ .

We use this representation because of the organization, clarity and insight it brings to the evaluation process. To apply it we need to address the following questions:

1. Is the summation over biases sufficiently complete?
2. Is σ determined by well-behaved processes?
3. Is ν_{ref} constant over the measurement interval?

Much of this paper is devoted to examining these questions.

1.1.2 Definitions

In the spirit of the recommendations put forward by the Consultative Committee for Time and Frequency (CCTF) [12], we make the following definitions for the uncertainty terms used in this paper.

By *Type A uncertainty* we mean how well successive frequency measurements agree with each other when the biases and reference oscillator are constant. It represents the instability of our measurements due to noise. It must be found by statistical analysis of numerous measurements. In (1) it is represented by σ .

By *Type B uncertainty* we mean how well we can determine the sum of biases. It is a combination of the uncertainties (also called Type B) for each bias in the summation. It includes statistical uncertainty in bias determination and estimates of biases not included in the summation. Type B uncertainty is not explicitly represented in (1).

By *combined standard uncertainty* (CSU) we mean the combination of Type A and Type B uncertainties. It represents how well the standard realizes the definition of the second.

By an *accuracy evaluation* we mean a series of measurements that determines the frequency difference between the reference oscillator and our primary standard, properly corrected for all known biases. The output of an evaluation is a calibration of the frequency of the reference oscillator ν_{ref} with a corresponding statement of the CSU of that calibration. The reference operates continuously, serving as a transfer standard to compare the primary standard with other local oscillators, to the NIST time scale AT1, and ultimately to the International Atomic Time (TAI) scale maintained

by the Bureau International des Poids et Mesures (BIPM).

In this paper uncertainties related to the frequency of the standard represent one standard deviation and are expressed in fractional frequency units.

1.1.3 Philosophy of uncertainty evaluation

The goal of an uncertainty evaluation is to not only calibrate a reference oscillator, but also to demonstrate how well that has been done, i.e. with what CSU. Referring to (1), we note first that the CSU cannot be less than the Type A uncertainty σ . For the CSU to approach σ , the Type B uncertainty in the determination of the biases must be smaller than σ . Therefore, we do not use frequency measurements under altered conditions to determine biases, since the Type A uncertainty of such additional measurements, also at least σ , would be compounded.

Instead, we determine biases by *leveraged* experiments. These are measurements of some auxiliary quantity, such as the frequency of another line in the spectrum. We then invoke a theoretical model to describe how the auxiliary quantity relates to the bias on the clock transition. In doing this we must predict how the uncertainty in the auxiliary measurement translates into uncertainty in the bias. We must also consider the validities of the models. We have now developed a complete set of leveraged experiments and models to determine all our known biases. These are discussed in Sections 4 and 5.

We also search for additional, unknown, or unexpected biases that might be present. This search includes *parametric measurements*, measurements of the standard's frequency with a known change in operating conditions. Examples are a change in the ambient magnetic field or in the microwave power used to excite the atoms. The search also includes some special techniques to find electronic biases. Such measurements form a check on the adequacy of our models and the completeness of the set $\delta\nu_j$. But we cannot guarantee that no significant biases remain. We believe the operators of a primary standard should consider every imaginable source of error and then look further for unanticipated biases. However, it is the nature of primary standards that there will always be subtle, undiscovered biases.

When one has exhausted the search techniques and arrived at a "best possible" estimate of uncertainty, one knows it is only a lower bound. What uncertainty should be reported to the BIPM? If one reports a larger uncertainty than computed from the uncertainties of known biases, one's frequency number will probably not be proven wrong. But if one reports the computed uncertainty, it is certainly optimistic. Which method would best serve the BIPM and the user community? We have no answer to this conundrum. In this paper we

report our computed Type B uncertainty and discuss its validity.

1.2 Outline of paper

Section 2 briefly reviews the physical description of NIST-7. Section 3 presents the theoretical models used for relating auxiliary measurements to biases. In Section 4 we discuss in detail how we determine the seven biases that we routinely include in the summation in (1). We also discuss the uncertainties associated with determining each bias. In Section 5 we discuss numerous additional sources of bias. In Section 6 we discuss the statistics of our measurements and the Type A uncertainty that we report. Section 7 describes the measurements that make up a typical evaluation. Finally, Section 8 summarizes the results and presents a combined standard uncertainty.

2. Description of the standard

NIST-7 is an optically pumped, thermal atomic-beam, microwave spectrometer designed specifically as the US primary frequency standard. Great care has been taken throughout its design and assembly to allow for the characterization and control of conditions that bias the measurement of the atomic resonance. Compared with the more conventional Stern-Gerlach type of magnetic state selection, optical pumping yields greater utilization of atomic beam flux. It produces a more symmetric Zeeman spectrum (Figure 6). It allows a more uniform magnetic field in the interrogation region. It provides greater spatial beam uniformity with better defined atomic trajectories while passing a wider distribution of atomic velocities. These aspects of the device are discussed later where they relate to accuracy. For completeness we begin with a brief outline of the physical system. Detailed descriptions of NIST-7 and discussion of its design can be found in [13-15].

2.1 Beam tube

A schematic of NIST-7, Figure 1, shows the atomic beam, optical pumping and detection regions, microwave excitation regions and magnetic shields (broken lines). The entire structure is symmetric about the midpoint. Separate regions of optical pumping and detection are located at each end. The components of the beam tube are briefly described below.

2.1.1 Microwave cavity

The Ramsey method of separated oscillatory fields [16] is used to interrogate the atoms. The cavity spans a drift region of length $L = 1.53$ m. The microwave interaction zones at the ends of the cavity are ring structures called De Marchi cavities [17]. These cavities have interaction zones of length $\ell = 2.3$ cm, across

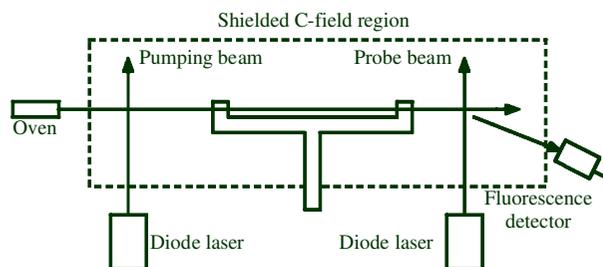


Figure 1. Schematic diagram of the NIST-7 beam tube. The atomic beam, optical pumping and detection regions, microwave excitation regions and magnetic shields (broken lines) are shown.

which the microwave magnetic field has a half-sine-wave profile. The atomic beam is defined by circular apertures of 3 mm diameter in the cavity ends.

The cavity is assembled from several parts (the cavity ends, the elbows, the straight sections, the central tee and the input coupling) joined with high-precision, flat, X-band flanges. During manufacture each part is carefully trimmed to a length corresponding to a half-integer multiple of the guide wavelength, in vacuum at the operating temperature of the standard. The uncertainty in the length of any part is approximately 10^{-6} m, the same as the alignment precision of the pinned flanges. This trimming operation places the flanges at deep nulls in the microwave field to minimize leakage from the joints. It also contributes to the tuning and end-to-end symmetry of the cavity. The cavity is quasi-isolated from the coaxial feed line with a 10 dB attenuator located right at the waveguide coupling pin. The loaded Q of the cavity is about 300.

2.1.2 Magnetic and electric fields

The requirements for both the microwave cavity orientation and the magnetic field uniformity are best served by a magnetic quantization field (C-field) parallel to the atomic beam. The C-field solenoid and shield package extend 2.34 m over the entire length of the beam tube and encompass the regions of optical state selection and state detection. The shielded solenoid without additional trim coils produces a magnetic field at the Ramsey cavities that differs from the mean field by a few parts in 10^4 at our nominal operating field of 5.64 μ T. Because of their long cylindrical geometry, the three nested magnetic shields produce a transverse shielding of order 10^6 but a longitudinal shielding of only order 10^3 . The high transverse shielding ensures that external perturbations can change only the axial field. The relatively poor axial shielding is compensated by active control of the current through the C-field coil. Degaussing coils are wound separately on each shield.

To control dc electric fields within the standard as well as to eliminate stray electric currents, all the substructures within the standard (microwave cavity, C-field solenoid, magnetic shields and vacuum system) are electrically isolated from one another.

Furthermore, no junctions of dissimilar metals are allowed. The substructures are intentionally made electrically common at one point. In this way, the entire beam tube is an electrostatic shield. At the same time, no unintentional currents can flow through the standard where they could generate undesirable magnetic fields.

2.1.3 Caesium ovens

The atomic beam emerges from caesium ovens located externally at each end of the beam tube. Each oven has a single-tube collimator, 3 mm in diameter and 10 cm long. The collimator produces an atomic beam with a half-angle divergence of 15 mrad. The beam is further collimated by the apertures of the microwave cavity. A typical atomic flux at the detection region is $5 \times 10^8 \text{ s}^{-1}$ when the oven temperature is 100 °C. For four years we used a hybrid, re-circulating collimator oven design [18]. We now use a simple, single-tube collimator followed by an absorptive aperture. This oven and its performance are also described in [18]. The oven temperature is now servo-controlled, but was not always controlled in the past. The external location, approximately 65 cm from the nearest cavity end, ensures that the heat of the oven will not create thermal gradients in the cavity that would distort end-to-end cavity phase measurements.

2.1.4 Laser light source

The optical pumping and detection light is generated by diode lasers whose wavelengths are stabilized using a saturated absorption feature in external caesium-vapour cells. The caesium transition is from the $6^2S_{1/2}$ ground state to the $6^2P_{3/2}$ excited state at 852 nm. The hyperfine structure of this transition is shown in Figure 2. The excited state hyperfine separations are from Tanner and Wieman [19]. The specific hyperfine transition used for state preparation is the $F = 4 \rightarrow F' = 3$, where F and F' indicate the ground and excited state hyperfine levels, respectively. Detection is accomplished by driving the $F = 4 \rightarrow F' = 5$ cycling transition. The two different optical frequencies were formerly generated from a single extended-cavity laser in conjunction with an acousto-optic modulator. Mirrors guided the beams from an optical table to the standard. But air turbulence caused fluctuations in the pointing of the laser beams. Recently, two, separate, distributed Bragg reflector (DBR) lasers have been used, set up on independent optical tables. Their light is transmitted to the standard by optical fibres. These changes have greatly reduced sensitivity to acoustic noise.

2.1.5 Optics

The beam tube has separate regions for optical state preparation and detection at each end of the beam tube to facilitate beam reversal. This also

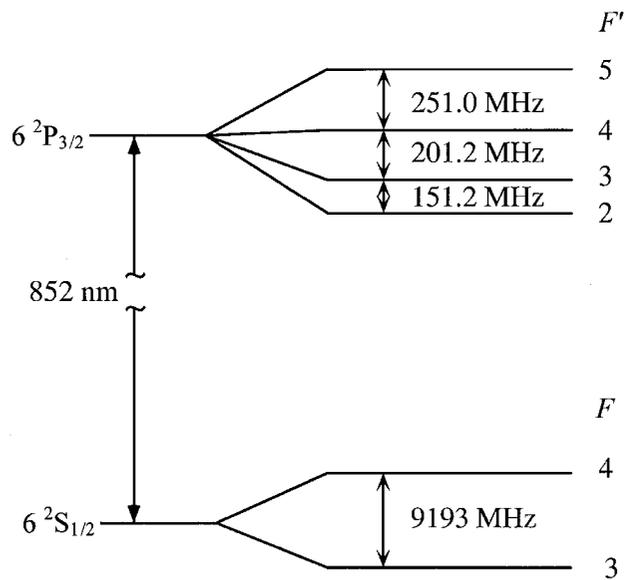


Figure 2. Hyperfine structure of the caesium optical transition.

permits simultaneous operation of counterpropagating atomic beams. The laser beams are injected into the system through normal incidence, anti-reflection coated windows. The laser beams are linearly polarized and perpendicular to the atomic beam. The state-preparation beam is retroreflected with its polarization rotated by $\pi/2$ rad to create scrambled polarization in the pumping region. This allows all atoms to be pumped without leaving any in “dark states” that would generate an unwanted background signal at the detector [20]. The fluorescence collection optics cover 2π steradians and incorporate spatial filtering to reject scattered light [13]. The fluorescence detector is located outside the vacuum system at the end of a glass light pipe. It consists of an unbiased, 5 mm^2 photodiode of low junction-capacitance, followed by a low-noise operational amplifier in a transconductance configuration. The transfer function is $1 \times 10^{10} \text{ V/A}$. The noise floor of the detector is $50 \mu\text{V}/\sqrt{\text{Hz}}$ and the effective bandwidth is 33 Hz.

2.1.6 Vacuum chamber

The main vacuum chamber is a cylinder 25 cm in diameter, 2.5 m long, and symmetric about its midpoint. It is pumped by two 20 l/s ion pumps. Our typical operating pressure is $3 \times 10^{-6} \text{ Pa}$, measured at the pumps. This is an order of magnitude lower than the pressure where we begin to see collisional attenuation of the atomic beam. The chamber is wrapped with a nonmagnetic, bifilar-wound heater plus thermal insulation. During operation, the chamber temperature is held constant at 39 °C.

2.2 Electronics

Numerous electronic systems are required to operate the standard and maintain its environment. The major ones are described below.

2.2.1 Microwave source

Figure 3 is a block diagram of the microwave synthesizer [21]. A 5 MHz quartz-crystal oscillator is phase-locked to the 5 MHz output of the reference oscillator, a hydrogen maser from the NIST timescale ensemble. This phase-locked loop ensures low phase noise at Fourier frequencies above 1 Hz while also preserving the long-term stability of the maser. To further reduce phase noise above 100 Hz, a 100 MHz quartz oscillator is phase-locked to the 5 MHz quartz oscillator. The 5 MHz quartz oscillator (multiplied by 8 to produce 40 MHz) replaces the internal frequency reference for a commercial direct digital synthesizer (DDS) [22]. The nominal output frequency of the DDS is 10 701 765 Hz. Potential spectral impurities in the DDS output are removed by a narrowband phase-lock to a 10.7 MHz crystal oscillator. The output of the 100 MHz oscillator is multiplied by 5 and mixed with the 10.7 MHz to produce 510 701 765 Hz. This signal is filtered, amplified and applied to a step recovery diode. The 18th harmonic at 9 192 631 770 Hz is then selected with a filter. This output frequency is ν_{ref} in (1) as it was synthesized from the reference oscillator. Using computer control we make small changes in the frequency of the DDS to tune the microwaves by known amounts with a resolution of 18 μHz . The tuning range is ± 180 kHz, sufficient for measuring the locations and lineshapes for all seven microwave transitions at nominal C-fields.

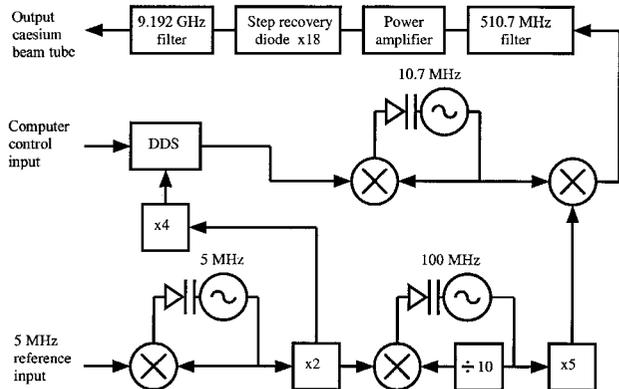


Figure 3. Block diagram of the microwave synthesizer. The 5 MHz input from the reference oscillator is multiplied to the caesium resonance frequency. Computer control of the direct digital synthesizer (DDS) permits tuning and modulation.

2.2.2 Frequency servo

Figure 4 is a block diagram of the digital frequency-control system [23]. The frequency servo performs slow square-wave frequency modulation of the microwave synthesizer by computer control of the DDS. The fluorescence signal from the caesium beam is detected by a photodiode, amplified, and digitized by an analogue-to-digital converter (ADC). A computer samples and demodulates the signal, inserting blanking intervals after each frequency change. The demodulation waveform is designed to produce deep nulls in the frequency response of the servo for multiples of the power-line frequency. The resulting error signal drives a software controller that steers the DDS centre frequency toward the centre of the atomic resonance, as described in Section 3.4. The frequency-steering corrections are stored by the computer as values of $\Delta\nu$ in (1). The mean frequency difference between the reference oscillator and the atomic resonance, together with its variance, is later calculated from these values. The frequency output of NIST-7 is thus not an electronic signal, but a table of frequency-difference numbers. Unlike previous frequency standards at the NIST, the frequency servo of NIST-7 does not lock a local oscillator to the caesium atomic resonance.

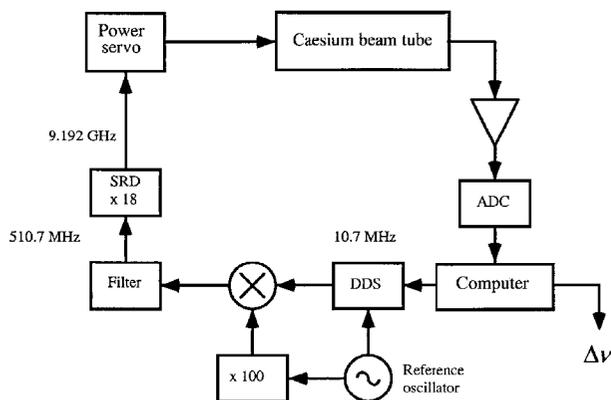


Figure 4. Block diagram of the frequency servo, showing the principal components used for frequency measurements.

2.2.3 Microwave power servo

A block diagram of the microwave power servo is shown in Figure 5. One half of the microwave power is sent to the beam tube. The other half is measured with a commercial radio-frequency (rf) power meter. A digital feedback control system monitors the power meter and adjusts an attenuator to maintain the desired power. The absolute accuracy of the commercial power sensor is not critical. Instead, we calibrate the power meter in terms of the atomic Rabi frequency. Once this calibration is performed, we can make changes in the microwave power with an uncertainty of 0.02 dB. All components of the microwave servo are enclosed in an rf-tight metal box to prevent microwave leakage

into the laboratory. All inputs and outputs to and from this box are configured to prevent microwave leakage without also introducing ground loops. Similar attention has been given to the components of the frequency synthesizer.

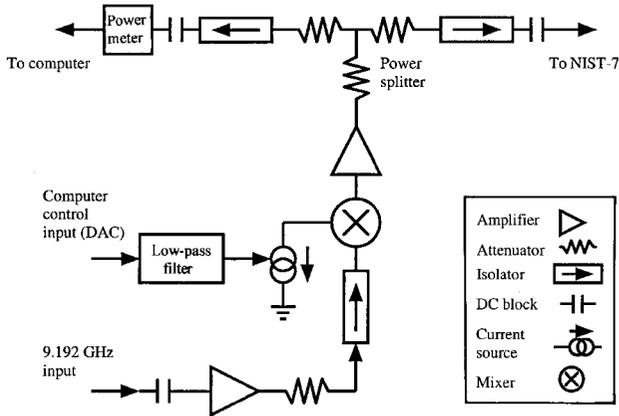


Figure 5. Block diagram of the microwave power servo. A mixer is used as a current-controlled attenuator.

2.2.4 C-field servo

The frequency servo periodically steps out from the clock transition to the field-sensitive $m = 1$ line and measures its position with a single modulation cycle. This information is used by the computer to make small corrections to the C-field current so that the position of the field-sensitive line is held constant.

2.3 Software

NIST-7 employs several computers for data acquisition, data analysis and implementation of digital servos. We have found that procedural programming languages such as C are not suitable for large applications. Instead, C++ was selected because it is an object-oriented language. C++ permits reuse of source code through the mechanism of software inheritance. Inheritance allows new software modules to contain all the tested functionality of any previously written modules without explicitly copying source code. Older features may be modified or removed, while new features may be added. Inheritance reduces the total amount of source code that must be tested, as only the new or modified features represent additional code. This also makes code maintenance simpler, as a particular algorithm or data structure resides in a single location, propagating through inheritance, not by source-code duplication. Four years of software development has produced TFLIB, a new library of software tools that exceeds 40 000 lines of C++ source code. Its object-oriented design has made it generally useful for time and frequency metrology. TFLIB has been used to build software applications not only for NIST-7, but also CRL-O1 [9], another optically pumped caesium thermal beam, and NIST-F1 [24], a caesium fountain.

2.4 Operational parameters

Table 1 summarizes the basic operational conditions for NIST-7, including a few not detailed elsewhere in this paper. The Rabi frequency and Ramsey linewidth correspond to the microwave excitation power used for normal operation, that is 2.5 dB below optimum.

Table 1. Typical operating parameters for NIST-7.

Parameter	Value
Oven temperature	100 °C
Mean atomic velocity	230 m/s
Width of velocity distribution	240 m/s
Interaction length ℓ	2.286 cm
Separation L of interaction regions	153.25 cm
Atomic beam aperture diameter	3 mm
Beam tube temperature	39 °C
Optical pumping transition	$F = 4 \rightarrow F' = 3$
Optical detection transition	$F = 4 \rightarrow F' = 5$
Magnetic field (C-field)	5.64 μ T
Zeeman line separation	39.5 kHz
Excitation Rabi frequency	1.8 kHz
Ramsey fringe linewidth	62 Hz

3. Theoretical background

As our evaluation technique is heavily model-dependent, we describe here the theoretical models we use in determining biases. Few of the results are new, but some of the formalism is not common to the clock literature. For reference we include a table of the many symbols used in the paper. A few of them have multiple meanings, but local definitions in different parts of the paper should avoid confusion.

A	autocorrelation of input noise
a	time integral of half Rabi frequency
B	cavity pulling coefficient
b	half Rabi frequency
C_x	wave-function amplitude
c	speed of light
D	denominator for most biases
E	electric field; error signal; end-to-end phase bias coefficient
F	angular momentum quantum number
f	frequency difference
f_L, f_R	aperture illumination fractions
$f(z)$	inhomogeneity function
G	ideal servo gain
g	servo gain; acceleration of gravity
g_J, g_I	g -factors
H	magnetic field
h	height above geoid
$h(\tau)$	function for l/L corrections to lineshapes
\hbar	Planck constant divided by 2π
I	atom flux
J_n	Bessel function of order n
$j(\tau)$	associated function for l/L corrections
K	ratio of signal to transition probability

L	length of drift region
l	length of excitation region
m	azimuthal quantum number
N	number of data points in a frequency measurement
n	index for frequency data; atomic density
P	transition probability
p	$\sqrt{b^2 + \Delta^2}$
Q	resonance line Q
q	residual fractional frequency difference
R	Ramsey fringe lineshape
R_k	unbiased frequency
r	servo convergence rate
S	signal
s_n	signal noise
T	transit time across drift region
t	time
U	evolution matrix
u	fractional frequency including end-to-end phase bias
v	velocity
W	temperature ratio in black-body formula
w	weights for frequency measurements
x	Ramsey dimensionless field parameter
Y	slope of background
z	longitudinal coordinate for inhomogeneity analysis
α, β	wave function amplitudes
Δ	half-angular frequency detuning
δ	prefix for frequency biases
ϵ	relative amplitude of inhomogeneity
θ	phase of microwave field
Λ	power and modulation coefficient for inhomogeneity
μ_B	Bohr magneton
ν	frequency
$\rho(T)$	transit-time distribution
σ	Type A frequency uncertainty; cross-section
$\sigma_y^2(\tau)$	Allan variance
τ	transit time across excitation region
τ_0	frequency servo cycle time
Φ	phase change across drift region
ϕ	end-to-end phase difference
ψ	Schrödinger wave function
ω	angular frequency: $\omega = 2\pi\nu$ whenever ω and ν have the same subscript

3.1 Zeeman spectrum

The hyperfine separation in caesium arises from the dipole interaction of the valence electron with the Cs^{133} nucleus. Because of the large nuclear spin, $I = 7/2$, the two hyperfine levels have angular momenta $F = 3$ and $F = 4$, with seven and nine magnetic sublevels, respectively. The Zeeman spectrum has seven σ transitions ($\Delta m = 0$) starting from each of the $F = 3$ sublevels. We use the initial (and final) m value to designate which σ -transition we are

referring to. Our frequency measurements are made on the $m = 0$ to $m = 0$ transition, often called the “clock” transition. There are also fourteen π -transitions ($\Delta m = \pm 1$) consisting of six overlapping pairs located in frequency between the σ -transitions, and two located just outside. The σ -transitions are excited by the microwave magnetic field component parallel to the static C-field, while the π -transitions are excited by microwave magnetic field components perpendicular to the C-field.

Figure 6 shows a Zeeman spectrum for NIST-7. The broad overlapping resonances are the Rabi pedestal lineshapes for the σ -transitions, while the narrow spikes are the unresolved Ramsey fringes. These lineshapes are fully described in Section 3.2. Pedestal overlap may cause a bias in the position of a Zeeman line. This effect, called Rabi pulling, is discussed in Section 5.2. Although the dominant microwave magnetic field component is oriented parallel to the C-field to excite the σ -transitions, the small dips between and outside the major resonances are indicative of weak excitation of the π -transitions for some atoms [25]. At higher C-fields and microwave power these transitions are clearly resolved.

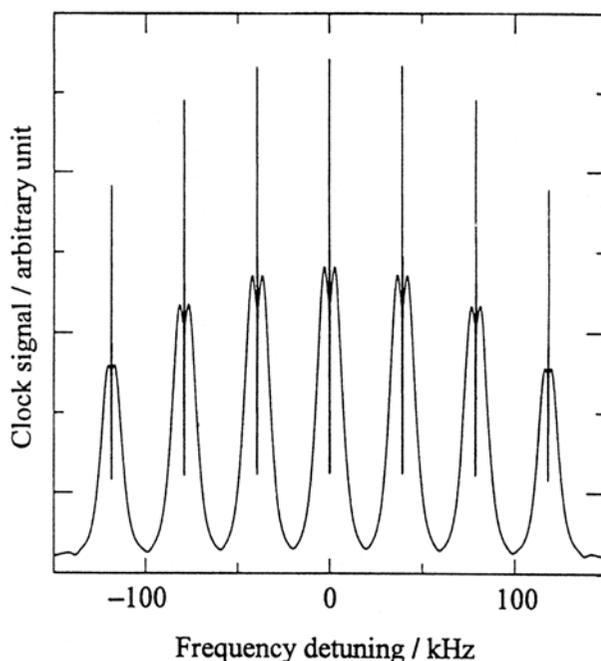


Figure 6. Zeeman spectrum for NIST-7. This spectrum was recorded at a C-field of $5.64 \mu\text{T}$ and optimum microwave power for the clock transition.

The magnetic field dependence of the σ hyperfine-transition frequencies is given by the Breit-Rabi formula [16, Eq. (III.120,121)]

$$\nu_m = \nu_{\text{hfs}}(1 + mx/2 + x^2)^{1/2}, \quad (2)$$

where m is the magnetic quantum number of the initial and final states,

$$x = (g_J - g_I)\mu_B H / (2\pi\hbar\nu_{\text{hfs}}), \quad (3)$$

g_J and g_I are the g -factors for the atom and nucleus, respectively, μ_B is the Bohr magneton, H is the magnetic field intensity, and $2\pi\hbar$ is Planck's constant. This relation is derived by applying the quantum mechanics of angular momenta to the dipole interactions of the atomic and nuclear magnetic moments with each other and with the C-field [16]. Corrections to the Breit-Rabi formula are insignificant unless the C-field is orders of magnitude larger than where we operate [26].

The relative size x of the first-order Zeeman shift is about 2×10^{-5} at the C-field we normally use. Hence a power-series approximation of the square root is adequate:

$$\nu_m = \nu_{\text{hfs}}[1 + mx/4 + x^2(1 - m^2/16)/2]. \quad (4)$$

For the field-dependent transitions the third-order term in x would be of order 1×10^{-15} . This is negligible compared with the observed fluctuations in the linear term. The clock transition has no first- or third-order contribution, while the fourth-order term is of order 2×10^{-20} .

If we define the Zeeman frequency ν_Z by

$$\nu_Z = (x/4)\nu_{\text{hfs}} = (g_J - g_I)\mu_B H/8\pi\hbar, \quad (5)$$

then (4) can be rewritten

$$\nu_m = \nu_{\text{hfs}} + m\nu_Z + 8\nu_Z^2(1 - m^2/16)/\nu_{\text{hfs}}, \quad (6)$$

where the magnetic field and g values no longer explicitly appear. Equation (6) is the model we use for the frequency of the σ hyperfine transitions.

3.2 Rabi and Ramsey lineshapes

In this section we describe the theory of the Rabi and Ramsey lineshapes observed in the Zeeman spectrum. The basic approximation is to treat the caesium atom as a two-level system, the ground state hyperfine levels. The small effects of other magnetic sublevels are discussed under Rabi and Ramsey pulling. The effects of optical transitions are discussed under black-body and light shifts.

As the C-field is nominally parallel to the atomic beam, the microwave cavity has an H-plane orientation. In this configuration the atoms experience a microwave field pulse of half-sine-wave profile as they pass through the cavity. Previous primary standards, such as NBS-6 and NBS-4, as well as most commercial standards, use a transverse C-field and an E-plane orientation of the cavity. The atoms then experience a microwave field pulse with a square profile. These different excitation pulse shapes cause differences in the lineshape, as noted below.

3.2.1 Quantum mechanics of a two-level system

We present here a formulation based on the time-dependent Schrödinger equation. The rudiments of this formulation have been published previously [27, 28]. An equally valid approach employing the density matrix is presented in the book by Vanier and Audoin [29].

The time-dependent Schrödinger equation for a two-level system excited by radiation at frequency ω can be written

$$\begin{aligned} i\hbar \frac{d}{dt} \psi &= \mathcal{H} \psi, \\ \psi &= \begin{pmatrix} C_a \\ C_b \end{pmatrix}, \\ \mathcal{H} &= \hbar \begin{pmatrix} \omega_a & 2b \cos(\omega t + \theta) \\ 2b \cos(\omega t + \theta) & \omega_b \end{pmatrix}. \end{aligned} \quad (7)$$

Here, ψ is the two-component wave function, \mathcal{H} is the Hamiltonian, and $\hbar\omega_a$ and $\hbar\omega_b$ are the energies of the upper and lower states, respectively. The interaction Rabi frequency is $2b = \mu_B g H_{\text{rf}}/\hbar$, where μ_B is the Bohr magneton, g is the Landé g factor, H_{rf} is the amplitude of the microwave magnetic field parallel to the C-field, \hbar is Planck's constant divided by 2π , and θ is the phase of the field at $t = 0$. The variables b and θ are real and may vary with time, but slowly compared with ωt .

The first step towards solving (7) is to use the rotating-wave approximation. We write the cosine in the interaction term as the sum of two exponentials. We then keep only the term that "rotates" in the same sense as the "precession" corresponding to the energy levels, where the terms in quotes have physical meaning for magnetic resonance of a spin-one-half system [30]. That is, we replace (7) by the rotating-field Hamiltonian

$$\mathcal{H}_{\text{rot}} = \hbar \begin{pmatrix} \omega_a & bc^{-i\omega t} e^{-i\theta} \\ be^{i\omega t} e^{i\theta} & \omega_b \end{pmatrix}, \quad (8)$$

where $\omega_a > \omega_b$. The principal correction caused by the "anti-rotating" part of the Hamiltonian is the Bloch-Siegert bias discussed in Section 5.4.

3.2.2 Phase-factoring transformation

The second step in solving (7) we call phase factoring. We define new probability amplitudes that differ from C_a and C_b by time-dependent phases:

$$\begin{aligned} \alpha(t) &= C_a(t) \exp[i\phi_a(t)], \\ \beta(t) &= C_b(t) \exp[i\phi_b(t)]. \end{aligned} \quad (9)$$

Probabilities are unaffected by such a phase change. The new amplitudes α and β are the components of a wave function ψ_p obeying

$$i\hbar \frac{d}{dt} \psi_p = \mathcal{H}_p \psi_p. \quad (10a)$$

By choosing the phases

$$\phi_a = (\omega + \omega_a + \omega_b)t/2 + \theta/2 \quad (10b)$$

$$\text{and } \phi_b = (-\omega + \omega_a + \omega_b)t/2 - \theta/2, \quad (10c)$$

the new Hamiltonian becomes

$$\mathcal{H}_p = \hbar \begin{pmatrix} -\Delta & b \\ b & \Delta \end{pmatrix}, \quad (10d)$$

where the coefficients b and Δ are both real. We have introduced the abbreviation $\Delta = (\lambda + d\theta/dt)/2$, where $\lambda = \omega - \omega_0$ is the detuning from the atomic resonance frequency $\omega_0 = \omega_a - \omega_b$. Either or both of b and Δ may be time-dependent. With these phase changes we have eliminated the rapid time-dependence and complex phase of the coupling coefficients. We have also made the new Hamiltonian real and traceless.

We further define α and β as those solutions of (10a) obeying the initial conditions $\alpha(0) = 1$ and $\beta(0) = 0$. The probability that a transition has occurred after excitation for a time τ is then represented by

$$P(\tau) = |\beta(\tau)|^2. \quad (11)$$

The form of (7) or (10a) guarantees that the normalization of the wave function is constant at all times:

$$|\alpha|^2 + |\beta|^2 = 1. \quad (12)$$

3.2.3 Rabi solutions

The general solution of (10) can be expressed by

$$\psi_p(t) = U(t, 0)\psi_p(0), \quad (13)$$

where the unitary evolution matrix $U(t, 0)$ satisfies (10a) and the initial condition that $U(0, 0)$ is the unit matrix $\mathbf{1}$. By the unitary requirement on a 2×2 matrix, the evolution matrix must have the form

$$U(t, 0) = \begin{pmatrix} \alpha(t) & -\beta^*(t) \\ \beta(t) & \alpha^*(t) \end{pmatrix}. \quad (14)$$

In general, $U(t, 0)$ cannot be found analytically when either b or Δ is time-dependent. We give below one exact solution and two approximate solutions.

Constant excitation

When b and Δ are both time-independent, (10) are coupled, linear, first-order differential equations with constant coefficients that can be solved by standard techniques, e.g. Laplace transforms. Here we use a matrix technique [27].

We write the evolution matrix as

$$U(t, 0) = \exp[-i\mathcal{H}_p t/\hbar]. \quad (15)$$

As a constant \mathcal{H}_p commutes with itself at all times, no time-ordering of operators is required in the

exponential. Expanding the exponential of the matrix gives

$$U(t, 0) = \cos(\mathcal{H}_p t/\hbar) - i \sin(\mathcal{H}_p t/\hbar) \\ = \cos(\mathcal{H}_p t/\hbar) - i[\sin(\mathcal{H}_p t/\hbar)/\mathcal{H}_p]\mathcal{H}_p. \quad (16)$$

Now the cosine is an even function of its argument, or a function of the square of its argument. From (10d), the square of the Hamiltonian is a constant times the unit matrix:

$$\mathcal{H}_p^2 = \hbar^2 p^2 \mathbf{1}, \quad (17)$$

where $p^2 = \Delta^2 + b^2$. Hence, $\cos(\mathcal{H}_p t/\hbar)$ is just $\cos(pt)$ times the unit matrix. Similarly the sine is an odd function of its argument, or the argument times an even function, as shown in the second line of (16). Substituting (17) into (16) we find the solution

$$U(t, 0) = \cos(pt)\mathbf{1} - (i/\hbar p) \sin(pt)\mathcal{H}_p. \quad (18)$$

The components are

$$\alpha(t) = \cos pt + i(\Delta/p) \sin pt,$$

and

$$\beta(t) = -i(b/p) \sin pt. \quad (19)$$

Note that β is pure imaginary.

The associated probability for an atom initially in state \mathbf{a} to make a transition to state \mathbf{b} is

$$P(t) = |\beta(t)|^2 \\ = (b/p)^2 \sin^2 pt, \quad (20)$$

which is often referred to as the Rabi formula [31]. It describes a lineshape with a strong central peak and decreasing side lobes separated by nulls. The height of the central peak reaches a maximum of unity when $bt = \pi/2$, or odd multiples thereof. We define optimum power to be the excitation power corresponding to the maximum transition probability at zero detuning.

Weak excitation

When b is weak, an approximate solution to (10d) valid to first order in b is easily found. First we assume b is zero and solve for α

$$\alpha(t) = \exp[i\Phi(t, 0)], \quad (21)$$

where

$$\Phi(t, 0) = \int_0^t \Delta(t') dt' \\ = [\lambda t + \theta(t) - \theta(0)]/2. \quad (22)$$

We then insert this result into the equation for β and obtain the approximate solution at time τ

$$\beta(\tau) = -i \exp[-i\Phi(\tau, 0)] \int_0^\tau b(t) \exp[2i\Phi(t, 0)] dt. \tag{23}$$

This solution is valid whenever the transition probability is small, i.e. in the wings of the Rabi lineshape, as well as when the excitation is weak.

If we define $b(t)$ as zero outside of the range 0 to τ , then the limits of integration in (23) can be extended to infinity and (23) can be recognized as a phase times the Fourier transform of $b(t)$. The transition probability becomes just the absolute square of the Fourier transform. The transform can be computed analytically for many forms $b(t)$ [32]. For b and θ constant in the range 0 to τ the transition probability becomes

$$P(\tau) = (2b/\lambda)^2 \sin^2(\lambda\tau/2), \tag{24a}$$

which is the same mathematical expression as obtained for single-slit diffraction. The λ^{-2} fall-off arises from the abrupt turn-on and turn-off of b .

When $b(t)$ has the shape of half a sine wave, we write

$$b(t) = (\pi b_0/2) \sin(\pi t/\tau), \quad 0 \leq t \leq \tau, \tag{24b}$$

where b_0 is not the peak, but the average height. If θ is constant and (24b) holds, the transition probability becomes

$$P(\tau) = \frac{(\pi^2 b_0/\tau)^2 \cos^2(\lambda\tau/2)}{(\lambda^2 - \pi^2/\tau^2)^2}. \tag{24c}$$

The λ^{-4} fall-off arises from the discontinuities in slope at 0 and τ . These approximations for large detuning are used to estimate Rabi pulling.

Small detuning

For the Ramsey solutions discussed in Section 3.2.4 we need results only for detuning that is small compared with b . To find an approximate solution for this case we first form the sum and difference of α and β in (10a):

$$i \frac{d}{dt}(\alpha \pm \beta) = \pm b(\alpha \pm \beta) - \Delta(\alpha \mp \beta). \tag{25}$$

For $\Delta = 0$ the solution of (25) is

$$\alpha \pm \beta = \exp[\mp ia(t)], \tag{26}$$

where the argument $a(t)$ is defined as the indefinite integral of $b(t)$:

$$a(t) = \int_0^t b(t') dt'. \tag{27}$$

A first-order solution is then obtained by substituting (26) into the detuning terms in (25) and integrating. Separating α and β we find

$$\alpha(t) = \cos a(t) + i \int_0^t \Delta(t') \cos[2a(t') - a(t)] dt' \tag{28a}$$

$$\beta(t) = -i \sin a(t) - \int_0^t \Delta(t') \sin[2a(t') - a(t)] dt'. \tag{28b}$$

If Δ is constant, (28) can be rewritten at time τ as

$$\alpha(\tau) = \cos a(\tau)[1 + i\Delta\tau h(\tau)] \tag{29a}$$

$$\beta(\tau) = -i \sin a(\tau)[1 - i\Delta\tau j(\tau)], \tag{29b}$$

where the new functions are

$$h(\tau) = (1/\tau) \int_0^\tau \cos[2a(t) - a(\tau)] dt \sec a(\tau) \tag{30a}$$

and

$$j(\tau) = (1/\tau) \int_0^\tau \sin[2a(t) - a(\tau)] dt \csc a(\tau). \tag{30b}$$

As the first-order corrections are out of phase with the zero-order solution, they are best interpreted as making small phase changes in α and β . They make only second-order corrections to the transition probability. These phase changes are used in the Ramsey fringe approximations in Section 3.2.4.

The integrals in (30) can be done analytically in two cases, which luckily are the cases of physical interest. When b is constant, we have $a(t) = bt$ and

$$h(\tau) = \tan(b\tau)/b\tau, \tag{31}$$

in agreement with the small-detuning approximation to (19). When $b(t)$ is a half-sine wave, we find from (24b)

$$a(t) = (b_0\tau/2)[1 - \cos(\pi t/\tau)] \tag{32}$$

and

$$a(\tau) = b_0\tau. \tag{33}$$

Then $h(\tau)$ contains the integral form of the zero-order Bessel function

$$h(\tau) = J_0(b_0\tau) \sec(b_0\tau). \tag{34}$$

In both cases $j(\tau)$ vanishes because $b(t)$ is symmetric about $\tau/2$.

If $b(t)$ were antisymmetric about $\tau/2$, then $a(t)$ would be zero. The transition probability would vanish at zero detuning. This case applies to the transverse component of the microwave field, which excites π -transitions when it has a half-cosine wave form, as it does for the cavity fields in NIST-7.

3.2.4 Ramsey solutions

General form and interpretation

In Ramsey excitation the atoms are partially excited in one region for a time τ_1 , allowed to drift freely for a time T , then excited again in a second region for a time τ_2 . We find the wave function for this case by successive applications of the evolution matrix:

$$\begin{aligned} \psi_p(\tau_2 + T + \tau_1) = \\ \mathbf{U}_2(\tau_2 + T + \tau_1, T + \tau_1) \mathbf{U}_{\text{free}}(T + \tau_1, \tau_1) \mathbf{U}_1(\tau_1, 0) \psi_p(0), \end{aligned} \quad (35)$$

where the subscripts 1 and 2 refer to the first and second excitation regions, respectively. \mathbf{U}_1 and \mathbf{U}_2 have the form given in (14). \mathbf{U}_{free} is the solution of (10a) when $\mathbf{b}(t)$ is zero. It is diagonal with elements given by (21) and (22), except for a change in the integration limits:

$$\Phi(T + \tau_1, \tau_1) = (\lambda T + \phi)/2, \quad (36)$$

where

$$\phi = \theta_2(T + \tau_1) - \theta_1(\tau_1). \quad (37)$$

The effect of \mathbf{U}_{free} is to alter the relative phases of α and β . The subscripts on θ allow for a difference in the microwave phases in the two excitation regions.

The matrix multiplications in (35) are easily carried out. The resulting transition amplitude is

$$\beta_2 \alpha_1 \exp(i\Phi) + \alpha_2^* \beta_1 \exp(-i\Phi). \quad (38)$$

Squaring, we obtain the transition probability

$$\begin{aligned} P(\tau_2 + T + \tau_1) = \\ |\alpha_2|^2 |\beta_1|^2 + |\beta_2|^2 |\alpha_1|^2 + 2\text{Re}\{\beta_2 \beta_1^* \alpha_2 \alpha_1 \exp(2i\Phi)\}. \end{aligned} \quad (39)$$

The first two terms constitute the Rabi-pedestal portion of the lineshape. They can be interpreted as the probability that a transition occurs in the first excitation region times the probability that it does not occur in the second excitation region, plus the probability that a transition occurs in the second excitation region times the probability that it does not occur in the first excitation region. The third term describes the Ramsey fringe resulting from the interference of the transition amplitudes in the two regions. When $T \gg \tau$ it describes rapid oscillations in detuning with an overall amplitude similar to that of the Rabi pedestal on which it sits.

Since $|\alpha|^2$ and $|\beta|^2$ are constrained by normalization, the Rabi pedestal and the Ramsey fringe are each constrained to the maximum value **0.5**. Optimum power, originally defined by $\mathbf{b} = \pi/2\tau$, now occurs at $\mathbf{b} = \pi/4\tau$ as there are two excitation regions.

Form for small detuning

In the vicinity of the central Ramsey interference fringe, when $T \gg \tau$ we can use the approximate solutions for small detuning (29) in (39). Keeping only first-order terms and assuming the two excitation regions are identical, we obtain

$$\begin{aligned} P_{\text{ped}}(2\tau) &= 2 \sin^2 a(\tau) \cos^2 a(\tau) \\ &= \sin^2 2a(\tau)/2 \end{aligned} \quad (40)$$

for the Rabi pedestal and

$$R(\lambda, T, \tau) = P_{\text{ped}}(2\tau) \cos[\lambda T + \lambda h(\tau)\tau + \phi] \quad (41)$$

for the Ramsey fringe. For single-velocity atoms, as in the caesium fountain, (40) and (41) provide an adequate description of the lineshape. As $\tau/T = \ell/L$, where ℓ and L are the lengths of the excitation and drift regions respectively, we refer to the term containing $h(\tau)$ as an ℓ/L correction. It will appear in the second-order Doppler, end-to-end phase, and inhomogeneity biases.

The full width of the Ramsey fringe, computed from (41) when the cosine vanishes, is $\Delta\nu_R = [2T + 2h(\tau)\tau]^{-1}$. This width actually decreases as microwave power and h increase. But this effect is unimportant when the lineshape is averaged over a broad velocity distribution.

3.2.5 Average over velocities

In the preceding analysis we have considered the atoms to have a constant velocity \mathbf{v} . In real atomic beams from thermal sources the atoms have a broad distribution of velocities $\rho_{\text{vel}}(\mathbf{v})$. This distribution is narrowed when magnetic state selection is used, but not when optical pumping is used. To average the preceding transition probabilities over atoms in the beam we introduce the transit-time distribution

$$\rho(T) = LT^{-2} \rho_{\text{vel}}(L/T), \quad (42)$$

where $\rho(T)$ and $\rho_{\text{vel}}(\mathbf{v})$ are both normalized to unit integral. We then write the averaged transition probability as

$$\langle P(\tau + T + \tau) \rangle \equiv \int_0^\infty \rho(T) P(\tau + T + \tau) dT. \quad (43)$$

Remember that $\tau = (\ell/L)T$ is also being integrated over.

As the width of the Ramsey fringes is proportional to \mathbf{v} , the velocity average superposes fringes of different widths. Like white-light fringes, only the central fringes survive this averaging. The rest “wash out”. The washing out occurs at smaller detuning as the transit-time distribution becomes broader. Figure 7 shows a typical Ramsey lineshape for NIST-7 recorded at our usual operating power, i.e. **2.5** dB below optimum

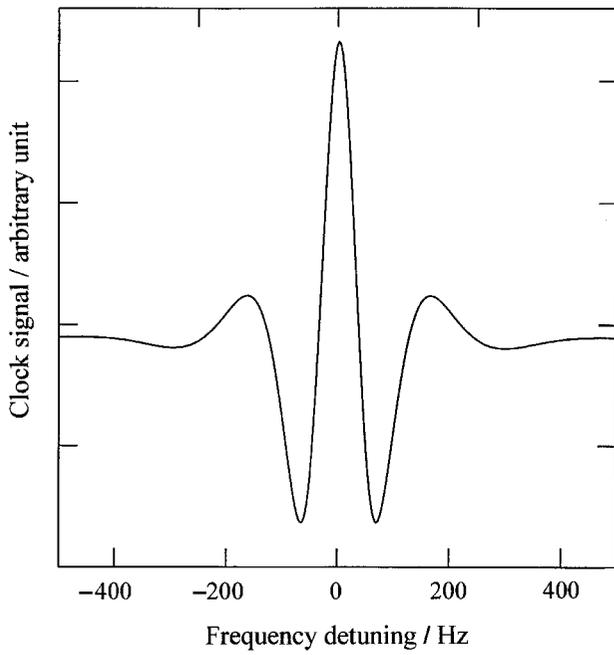


Figure 7. Ramsey lineshape for NIST-7 at 2.5 dB below optimum power.

of the transit-time distribution times the transition probability. This result is the basis for our method of determining velocity distributions [33, 34].

3.3 Bias models for slow square-wave modulation

Several biases depend on the transit time of an atom through the beam tube. These biases must then be averaged over a distribution of transit times with appropriate weighting factors to reflect their relative contributions to the signal. Other biases depend explicitly on the resonance lineshape. Both depend on the method used to find the line centre. In this section we describe these dependencies explicitly.

Let the dc clock signal at a given detuning λ be denoted by $S(\lambda)$. The servo measures signals $S_{\pm}(\lambda) \equiv S(\lambda \pm \omega_{\text{mod}})$, where ω_{mod} is the modulation amplitude in angular frequency units. It then adjusts λ to make $S_{+}(\lambda) = S_{-}(\lambda)$. The signal is proportional to the average transition probability defined by (43).

$$S(\lambda) = K \langle P(\lambda, \tau + T + \tau) \rangle, \quad (44)$$

where the coefficient K includes the beam flux, detection efficiency, preamplifier gain, analogue-to-digital conversion, etc.

3.3.1 Biases attributable to phase changes

As λ is small when the servo locks to the Ramsey fringe, we can use the approximate forms (40) and (41) in the signal. Then

$$S_{\pm}(\lambda) = K \left\langle \frac{1}{2} \sin^2 2a(\tau) \{1 + \cos[(\lambda \pm \omega_{\text{mod}})(T + h\tau) + \phi]\} \right\rangle. \quad (45)$$

When both λT and ϕ are very small the cosine can be approximated, so the servo resonance condition becomes

$$S_{+}(\lambda_{\tau}) - S_{-}(\lambda_{\tau}) = -K \langle \sin^2 2a(\tau) [\lambda_{\tau}(T + h\tau) + \phi] \times \sin \omega_{\text{mod}}(T + h\tau) \rangle \equiv 0. \quad (46)$$

Solving (46) for the resonance detuning λ_{τ} , we find

$$\lambda_{\tau} = -\langle \phi \sin^2 2a(\tau) \sin[\omega_{\text{mod}}(T + h\tau)] \rangle / D, \quad (47a)$$

with

$$D = \langle (T + h\tau) \sin^2 2a(\tau) \sin[\omega_{\text{mod}}(T + h\tau)] \rangle. \quad (47b)$$

The denominator D is just twice the slope of the Ramsey fringe (41) evaluated at $\lambda = -\omega_{\text{mod}}$. Equation (47) is used for the second-order Doppler and end-to-end phase biases treated in Sections 4.2 and 4.4.

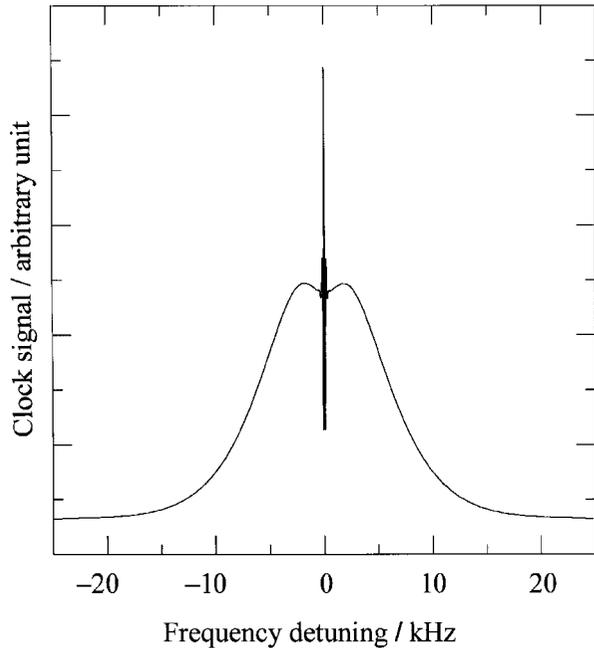


Figure 8. Rabi pedestal lineshape for NIST-7 at 2.5 dB below optimum power.

power. Only one side lobe on each side of the central fringe is distinct.

Figure 8 shows a typical Rabi pedestal for NIST-7 taken at a microwave power 2.5 dB below optimum. The side lobes have been totally washed out. The spike in the centre is the unresolved Ramsey fringe. The dip in the centre arises from slow atoms that are excited above their individual optimum powers.

Note that when (41) is substituted into (43) we have an expression close to the Fourier cosine transform

For the broad transit-time distributions found in NIST-7, the factor $T+h\tau$ in the denominator effectively shifts $\rho(T)$ to higher T values. This causes the integral ratio in (47) to depend appreciably on both microwave power and modulation amplitude. The modulation amplitude dependence means that the dc lineshape is asymmetric [35, 36]. For NIST-7 this asymmetry is very small, only 10^{-4} to 10^{-5} of the linewidth, but still significant for calculated biases.

3.3.2 Biases attributable to background variation

Another cause of frequency biases is a change in the background signal or the microwave field amplitude with detuning. If a background signal is present with a linear variation $KY\lambda$, then the signal difference in (46) is augmented by $2KY\omega_{\text{mod}}$. This leads to a background bias

$$\lambda_{\text{b}} = 2Y\omega_{\text{mod}}/D. \quad (48)$$

This result is used in the discussion of Rabi pulling in Section 5.2.

3.3.3 Biases attributable to resonance inhomogeneity

Suppose the atomic resonance frequency ω_0 varies along the beam path. This happens if the magnetic field varies (Section 5.1), or if changes are induced by the microwave field in or near the excitation regions. For the moving atoms in the beam such changes vary in time. For Ramsey excitation the accumulated phase in the drift region becomes

$$\Phi(T + \tau, \tau) = \frac{1}{2} \left[\omega T - \int_{\tau}^{T+\tau} \omega_0(t) dt + \phi \right] \quad (49)$$

in place of (36). This reduces to (36) if we define λ as $\omega - \bar{\omega}_0$, where $\bar{\omega}_0$ is the average resonance frequency in the drift region. With this refinement the preceding analysis remains valid.

However, biases do occur when the resonance frequency ω_i in the excitation region i ($i = 1$ or 2) differs from its average in the drift region by $\Delta\omega_i = \omega_i - \bar{\omega}_0$. Then the Δ in (28) is replaced by $(\lambda - \Delta\omega_i)/2$. This replacement causes a change in the phase of α . From (39) this change in phase augments the cosine argument in (41) by $-\Delta\omega h\tau$, where

$$\overline{\Delta\omega} = (\overline{\Delta\omega_1} + \overline{\Delta\omega_2})/2 \quad (50)$$

is the average offset of the resonance frequency in the two excitation regions. Equation (50) is also the offset of the Rabi pedestal, since the first two terms in (39) measure the frequency only in the excitation regions, averaging the offset for the two regions.

When this additional phase is inserted in place of ϕ in (47), we obtain the inhomogeneity shift

$$\lambda_{\text{in}} = \overline{\Delta\omega}(\ell/L)A, \quad (51)$$

where

$$A = \langle hT \sin^2 2\alpha(\tau) \sin \omega_{\text{mod}}(T + h\tau) \rangle / D. \quad (52)$$

The factor A , modelled after those used by De Marchi et al. [37], is of order unity and contains all the dependence on microwave power and modulation amplitude. It is used in the discussion of magnetic field inhomogeneity in Section 5.1.

3.4 Servo theory

The slow square-wave frequency servo used with NIST-7 makes an initial estimate f_0 for the resonance frequency. It then measures the signal level at a fixed offset ν_{mod} , the modulation amplitude, on either side of the initial estimate. The difference between the two signal levels is the error signal. From it we construct a new estimate for the resonance frequency and the process is repeated. The set of frequency estimates so obtained becomes the raw data of frequency measurements by NIST-7.

Let f_n be the n -th estimate for the microwave frequency. The error signal associated with this detuning is

$$E_n = S_-(f_n) + s_- - S_+(f_n) + s_+, \quad (53)$$

where $S_{\pm}(f)$ is the signal level on either side of the line as given by (45) and s_{\pm} is the signal noise accompanying the measurement of S_{\pm} . For f_n close to the resonance frequency f_{res} , we use a Taylor expansion about f_{res} to write

$$E_n = S_-(f_{\text{res}}) - S_+(f_{\text{res}}) + [S'_-(f_{\text{res}}) - S'_+(f_{\text{res}})](f_n - f_{\text{res}}) + s_n, \quad (54)$$

where S' is the derivative of S with respect to f , and $s_n = s_- - s_+$. By definition of f_{res} the first two terms cancel. We define the difference between derivatives in brackets to be the reciprocal of the ideal gain G . Then (54) reduces to

$$E_n = G^{-1}(f_n - f_{\text{res}}) + s_n. \quad (55)$$

A frequency correction is computed by multiplying the error signal by a gain factor g . This correction is then subtracted from the previous frequency estimate to obtain the next frequency estimate

$$f_{n+1} = f_n - gE_n. \quad (56)$$

In terms of the initial estimate, the $(n+1)$ -th estimate becomes

$$f_{n+1} = f_0 - g \sum_{j=0}^n E_j. \quad (57)$$

The summation of error signals is the digital counterpart of an integrator in an analogue servo. When we substitute (55) into (56) we obtain the recursion relation

$$f_{n+1} - f_{res} = (1 - g/G)(f_n - f_{res}) - g s_n. \quad (58)$$

This is the basic recursion relation for our servo theory.

If we further define the convergence rate $r = 1 - g/G$, then the solution of (58) can be written

$$f_n - f_{res} = r^n(f_0 - f_{res}) - g(s_{n-1} + r s_{n-2} + r^2 s_{n-3} + \dots + r^{n-1} s_0). \quad (59)$$

If the absolute value of r is less than 1, then (59) shows that the error in the initial estimate is reduced, as is the effect of noise many iterations ago. Table 2 shows how the value of g , or r , determines the behaviour of the servo. As the slope of the lineshape can be measured, the ideal gain can be determined and the actual gain set close to it for rapid convergence.

Table 2. Servo behaviour versus convergence rate.

Gain	Convergence rate	Servo behaviour
$g < 0$	$r > 1$	Diverges monotonically
$0 < g < G$	$0 < r < 1$	Converges monotonically
$g = G$	$r = 0$	Converges immediately
$G < g < 2G$	$-1 < r < 0$	Converges with oscillation
$g > 2G$	$r < -1$	Diverges with oscillation

The effect of the signal noise s_n is to introduce noise on to the frequency record. If the noise is white, as from atomic-beam shot noise, then it has zero mean and is uncorrelated with itself at another time:

$$\{s_n\} = 0 \quad (60)$$

$$\{s_m s_n\} = 2A \delta_{mn}, \quad (61)$$

where the curly braces denote an average over a large number of samples and A is a measure of the signal noise on one side of the line. Then the mean of the frequency record $\bar{f} = \{f_n\}$ is just f_{res} with no bias. But if the mean of the noise is $\bar{s} \neq 0$, then the frequency record acquires a bias. Summing the power series in r we obtain from (59) the mean frequency

$$\bar{f} = f_{res} - g\bar{s}/(1 - r). \quad (62)$$

Such a bias could occur if, for example, the detection laser's amplitude or the signal amplifier's gain varied in synchrony with the modulation.

The statistics of the frequency record reflect both the statistics of the noise and the convergence rate of the servo. Applying (61) to the square of (59) we obtain the variance

$$\{(f_n - f_{res})^2\} = 2g^2 A/(1 - r^2) \quad (63)$$

$$= 2G^2 A(1 - r)/(1 + r), \quad (64)$$

where the second form takes into account the relationship between g and r . The variance is reduced by using a gain less than the ideal gain (less noise is written on to the frequency record), but convergence is slower. This effect is readily observable when error signals and frequencies are plotted together in real time.

We also see from (59) that use of a non-ideal gain introduces correlation between successive values of frequency, even when no correlation exists in the input signal noise. We can compute this correlation in the same way that we obtained the variance:

$$\{(f_{n+k} - f_{res})(f_n - f_{res})\} = r^k \{(f_n - f_{res})^2\}, \quad (65)$$

valid for $k \geq 0$. Thus the correlation decreases by the convergence factor for each cycle of separation.

A common measure of frequency stability is the two-sample or Allan variance. Its basic definition is [38]

$$\sigma_y^2(k\tau_0) = \{(\bar{y}_{i+k} - \bar{y}_i)^2\}/2, \quad (66)$$

where \bar{y}_i is the average of the relative frequencies y_i to y_{i+k-1} and τ_0 is the cycle time. For $k = 1$ the Allan variance becomes

$$\sigma_y^2(\tau_0) = \{y_i^2\} - \{y_{i+1} y_i\} \quad (67)$$

$$= (1 - r)\{(f_i - f_{res})^2\}/\nu_{\text{hfs}}^2. \quad (68)$$

We can compute the Allan variance for small values of k in a similar fashion. For example,

$$\sigma_y^2(2\tau_0) = (1 + r)(1 + r/2)\sigma_y^2(\tau_0)/2. \quad (69)$$

For large values of k the Allan variance is dominated by the positive terms in the square of the averages. Using (64), (65) and (66) we find

$$\sigma_y^2(k\tau_0) = 2G^2 A/k\nu_{\text{hfs}}^2, \quad (70)$$

with corrections smaller by the factor $1/k$. This result is independent of the servo gain. It describes white frequency noise, the type expected from white input noise. Extrapolating (70) back to small values of k , we see that the results (68) and (69) are smaller if the gain is less than ideal and larger if it is greater. By using the ideal gain we obtain (70) for all values of k .

We apply these results in the noise discussion of Section 6.

4. Frequency biases for which corrections are made

There are seven frequency biases affecting NIST-7 for which corrections are routinely made. We write the sum of these biases in (1) as

$$\sum_j \delta\nu_j = \delta\nu_{\text{QZ}} + \delta\nu_{\text{D}} + \delta\nu_{\text{C}} + \delta\nu_{\text{E}} + \delta\nu_{\text{DC}} + \delta\nu_{\text{B}} + \delta\nu_{\text{G}}. \quad (71)$$

The $\delta\nu$ s refer to the quadratic Zeeman, second-order Doppler, cavity pulling, end-to-end phase, distributed cavity phase, black-body, and gravitational biases, respectively. The first four depend on operating conditions that may change from one evaluation to the next, or even from one frequency measurement to the next. The last three, which involve the atomic beam's alignment, the interior temperature of the beam tube, and its physical location, are stable.

Each bias is determined with the aid of information from *leveraged* experiments. Theoretical models are then applied to relate the results of these experiments to the determination of the biases. This procedure determines each bias with an uncertainty (Type B) smaller than the Type A uncertainty of a frequency measurement. The details for each bias are discussed below.

4.1 Quadratic Zeeman effect

This bias arises from the magnetic field H (C-field) deliberately imposed on the atoms. It is the largest bias, hence it must be measured with the smallest relative uncertainty. Our leveraged experiment is measurement of the frequency of an adjacent microwave transition that depends much more strongly on the magnetic field.

Our model, the Breit-Rabi formula, was presented in (6). From it we find the quadratic Zeeman bias for the clock transition is simply

$$\nu_0 - \nu_{\text{hfs}} = \delta\nu_{\text{QZ}} \equiv 8\nu_Z^2/\nu_{\text{hfs}}. \quad (72)$$

This shift is always positive and is independent of microwave power and modulation amplitude.

To find ν_Z we measure the frequency of the adjacent field-dependent $m = 1$ transition in the Zeeman spectrum

$$\nu_1 - \nu_{\text{hfs}} = \nu_Z + 7.5\nu_Z^2/\nu_{\text{hfs}}. \quad (73)$$

The difference between the frequencies of the $m = 0$ and $m = 1$ lines is just

$$\nu_1 - \nu_0 = \nu_Z - \nu_Z^2/2\nu_{\text{hfs}}. \quad (74)$$

Solving (74) iteratively we obtain

$$\nu_Z = \nu_1 - \nu_0 + (\nu_1 - \nu_0)^2/2\nu_{\text{hfs}} \quad (75)$$

with negligible error. The value of ν_1 is maintained by the C-field servo at a fixed value by adjusting the C-field current as needed. The value of ν_0 is the result of measurement of the clock transition. These numbers are used in (75) and (72) to compute $\delta\nu_{\text{QZ}}$.

The accuracy of the value of ν_Z obtained from (75) is limited by the ability of the C-field servo to hold ν_1 constant in the presence of fluctuations in the C-field arising from fluctuations in the external field. From Allan variance measurements of ν_1 with the servo active, the uncertainty in ν_1 is of the order of 0.01 Hz. This corresponds to an uncertainty of 0.01 Hz in ν_Z as

computed by (75). When ν_Z is about 40 kHz, where we normally operate, an uncertainty of 0.01 Hz in ν_Z corresponds to an uncertainty in $\delta\nu_{\text{QZ}}$ of 10^{-6} Hz or $10^{-16}\nu_{\text{hfs}}$ as computed by (72). This is the uncertainty we assign to this bias.

As measured, both ν_1 and ν_0 contain other biases. These biases are nearly identical for the two transitions, but do differ slightly because the $m = 1$ transition has a slightly smaller matrix element and is thus not quite as strongly excited by the same microwave field strength. By subtracting the two frequencies we are left with residual differences in these other biases of a few times $10^{-15}\nu_{\text{hfs}}$. We could compute these small differences, but they are much less than the fluctuations of 0.01 Hz in ν_1 .

To check our model for the quadratic Zeeman bias we operated NIST-7 at five different C-field values, corresponding to the ν_Z/kHz values 19.7, 26.3, 32.9, 39.5 and 52.7. The quadratic Zeeman bias varied from 0.34 Hz to 2.41 Hz, a span of 2×10^{-10} in fractional frequency. But after subtracting this and other known (essentially constant) biases, we found that the corrected frequencies had a scatter of only $3 \times 10^{-15}\nu_{\text{hfs}}$, equal to the Type A frequency uncertainty of the measurements at each C-field value.

4.2 Second-order Doppler effect

This bias arises from the special theory of relativity. For a single atom with longitudinal velocity v much less than the velocity of light c , it is given by

$$\delta\nu_{\text{D}} = -\frac{1}{2}\frac{v^2}{c^2}\nu_{\text{hfs}}. \quad (76)$$

For NIST-7, the mean velocity is about 230 m/s, making the second-order Doppler shift about $-3 \times 10^{-13}\nu_{\text{hfs}}$. This bias is always negative.

As the atomic beam includes atoms with a range of velocities, v^2 must be averaged over the velocity distribution with weighting factors for transition probability and modulation amplitude. We can do this by adapting the theory presented in Section 3.3.1. To a transversely moving atom the microwave frequency ω appears to be shifted to $\omega(1 + v^2/2c^2)$. Then the detuning λ should be replaced by $\lambda + (v^2/2c^2)\omega$. We make this substitution in the argument of the cosine in (45), and define the new terms to be the time-dependent phase ϕ_{D} . Recalling that $v = L/T$, we find

$$\phi_{\text{D}} = (L^2/2c^2)(\omega/T)(1 + h\ell/L). \quad (77)$$

This phase is very small, so it can be used in (47) in place of ϕ . The final result for the second-order Doppler shift is

$$\lambda_{\text{D}} = -(L^2\omega_0/2c^2) \times \langle T^{-1}(1 + h\ell/L) \sin^2 2a(\tau) \sin \omega_{\text{mod}}(T + h\tau) \rangle / D, \quad (78)$$

where D was defined in (47b). This relation is our model for this bias.

Neglecting $h\tau$ relative to T (78) simplifies to

$$\lambda_D = -(L^2\omega_0/2c^2)\langle T^{-1} \sin^2 2a(\tau) \sin \omega_{\text{mod}}T \rangle / D_0, \quad (79)$$

where D_0 is D without the $h\tau$ terms. This approximation is low by roughly $\ell/2L$. There is no singularity at small T because both the transit-time distribution and the modulation factor vanish there. The ℓ/L corrections included in (78) but not in (79) add about 0.5 % to the bias, so they are retained in our model.

The primary limitation to the accuracy of the second-order Doppler bias is in the determination of the velocity distribution. The leveraged experiment is the measurement of Ramsey lineshapes at several power levels. These lineshapes are then Fourier-transformed and combined in a weighted manner described in [33, 34] to obtain the velocity distribution. The method is based on the same lineshape theory derived in Section 3.2. The method also yields a value for b_0 to use in (78) or (79). Velocity distributions obtained by a time-of-flight method using pulsed optical pumping agree well with those from the Fourier transform method [39].

Repeatability of the computation of the second-order Doppler bias from a velocity distribution was investigated using experimental data. One hundred sets of Ramsey lineshape data were taken consecutively. A velocity distribution was extracted from each data set. The scatter in the Doppler biases computed from these velocity distributions was 0.06 %, corresponding to an uncertainty of 2×10^{-16} in fractional frequency of the standard. This scatter comes primarily from noise in the Ramsey lineshape measurements. Larger scatter, up to 0.5 % of the Doppler bias, has been observed in the past when the oven temperature was not well controlled.

Systematic error in the Fourier-transform method was tested by generating theoretical Ramsey lineshapes from (41) with an experimental velocity distribution. A reconstructed velocity distribution was then extracted from these lineshapes [34]. The second-order Doppler bias computed from the original and the reconstituted velocity distributions differed by 0.2 % or 6×10^{-16} in fractional frequency.

At a power level 2.5 dB below optimum, a change of 5 % in the Rabi frequency $2b$, corresponding to a change of 0.4 dB in microwave power, causes the second-order Doppler bias to change by 3.5 %, or $10^{-14}\nu_{\text{hfs}}$. Our power servo maintains the power constant to 0.02 dB, reducing the uncertainty in the second-order Doppler correction to $5 \times 10^{-16}\nu_{\text{hfs}}$.

Considering all these uncertainties, we assign a Type B uncertainty of $1.0 \times 10^{-15}\nu_{\text{hfs}}$ to the second-order Doppler bias.

There is no first-order Doppler effect because the exciting radiation is a standing wave and the beam atoms are moving transversely to it. Any residual imperfections in the standing wave or atom trajectories

can be described by small phase variations along an atom's path. The effects of these phase variations are discussed in Section 4.5.

4.3 Cavity pulling

This bias arises from the variation of the microwave field amplitude with microwave frequency due to the cavity resonance or other characteristic of the microwave circuit. The change in microwave field amplitude causes a change in the signal level which the frequency servo misinterprets as a tuning error. The phenomenon is most easily described by partial derivatives of the signal [40]. The relation is derived as follows.

Near resonance, the measured signal on the two sides of the lineshape can be approximated by

$$S_{\pm} = S(b, \pm\omega_{\text{mod}}) + \frac{\partial S(b, \pm\omega_{\text{mod}})}{\partial b} \delta b_{\pm} + \frac{\partial S(b, \pm\omega_{\text{mod}})}{\partial \lambda} \lambda. \quad (80)$$

To the accuracy needed the signal is a symmetric function of detuning. Hence, the b derivative is symmetric and the λ derivative is antisymmetric in λ . The change in b arises because b depends on detuning. This dependence is on a broader scale than the width of the Ramsey fringe or the Rabi pedestal, so we can write

$$\delta b_{\pm} = \pm (db/d\lambda) \omega_{\text{mod}}. \quad (81)$$

The resonance condition $S_+ = S_-$ then requires the detuning

$$\lambda_C = -\frac{\partial S/\partial b}{\partial S/\partial \lambda} \times \frac{db}{d\lambda} \omega_{\text{mod}}, \quad (82)$$

where the partial derivatives are evaluated at b, ω_{mod} . The signal derivatives are just K times the corresponding lineshape derivatives. Our final expression for the cavity pulling bias is then

$$\delta \nu_C = -\frac{\partial P/\partial b}{\partial P/\partial \lambda} \times \frac{db}{d\lambda} \times \nu_{\text{mod}}. \quad (83)$$

The detuning slope $db/d\lambda$ is a property of the microwave circuit, independent of the atoms. For a cavity detuning $\delta \omega_C$ much less than the cavity linewidth $\Delta \omega_C$, it can be approximated by

$$db/d\lambda = b_c \delta \omega_C / \Delta \omega_C^2, \quad (84)$$

where b_c is the value of b at the cavity resonance.

For the Ramsey fringe the b derivative can be computed from (40) and (41). The λ derivative is one half of (47b). Neglecting $h\tau$ compared with T , the cavity pulling bias becomes

$$\lambda_C = \langle 2\tau \sin 4b_0\tau (1 + \cos \omega_{\text{mod}}T) \rangle (db/d\lambda) \omega_{\text{mod}} / D_0. \quad (85)$$

This bias may be positive or negative depending on the microwave power level and detuning slope. It vanishes when the cavity is perfectly tuned, or when the microwave power is adjusted near optimum. It increases with modulation amplitude.

Equation (82) applies equally to the Rabi pedestal and the Ramsey fringe. Because the cavity resonance is much broader (megahertz) than the Rabi pedestal, we expect $db/d\lambda$ to be nearly the same for the pedestal as for the fringe. Thus, if we can measure the shift of the pedestal and compute the partial derivatives numerically (or measure them), we can infer $db/d\lambda$ and the shift of the Ramsey fringe.

For the pedestal the modulation amplitude ω_{mod} is about $2L/\ell$ times larger than for the fringe, while the slope of the lineshape $\partial S/\partial\lambda$ is about $2L/\ell$ times smaller. Consequently the cavity-pulling shift of the pedestal is about $4L^2/\ell^2$ times larger than the corresponding shift of the fringe. Although we cannot measure the pedestal position as well, by the $2L/\ell$ ratio of slopes, we still gain one factor of $2L/\ell$. This factor is our leverage.

Our procedure is to measure the offset of the pedestal from the centre of the Ramsey fringe. We obtain partial derivatives of the pedestal lineshape from numerical evaluation of the lineshape model in Section 3.2, or from experimental measurements. We then solve (82) for the detuning slope $db/d\lambda$. Finally, we evaluate (85) numerically or (82) with experimental fringe data to obtain the cavity pulling of the fringe.

The measured pedestal offset of the NIST-7 clock transition is about 2 Hz with 10 % uncertainty for a measurement time of 1000 s. The corresponding detuning slope is about $0.06b/\text{MHz}$. The computed fringe shift $\delta\nu_{\text{C}}$ at typical operating conditions is $6 \times 10^{-15}\nu_{\text{hfs}}$ with an uncertainty of 10 %.

4.4 End-to-end cavity phase bias

This bias arises from a difference in the phase of the microwave field in the two excitation regions. If this phase difference is ϕ , the bias is given by $\delta\nu_{\text{E}} = -\phi/2\pi T$, where T is the transit time of an atom between the two excitation regions. As with the second-order Doppler shift, T must be averaged over the velocity distribution with weighting factors for transition probability and modulation amplitude. The result is the ratio of transit-time integrals presented in (47). This relation is our model for this bias. For brevity we write the integral ratio as the coefficient E times $\omega_0\phi$ so that $E\phi$ is in fractional frequency units. A typical value for E is $2.4 \times 10^{-9}/\text{rad}$.

The atoms see the end-to-end phase difference change sign when the direction of the atomic beam is reversed. If atom trajectories in the two beam directions are equivalent, the magnitude of the phase difference should remain the same. Then frequency measurements made in both beam directions can be

combined to measure ϕ , or to remove it from the evaluated frequency. For one measurement in each direction the algebra is as follows [6]. Let u be the fractional frequency difference between NIST-7 and the reference with the biases discussed in Sections 4.1, 4.2 and 4.3 removed. Then for one beam direction we define the residual fractional frequency difference q with the end-to-end phase bias also removed:

$$q = u - E\phi \pm \sigma. \quad (86)$$

For the opposite beam direction we define a corresponding residual fractional frequency difference q' :

$$q' = u' + E'\phi \pm \sigma'. \quad (87)$$

The biases removed to obtain u' may differ from those removed to obtain u , as the velocity distributions often differ slightly in the two beam directions. Similarly, E' may differ from E by a few parts in 10^2 . The Type A uncertainty σ' should be uncorrelated with σ .

If no other biases that change with beam direction (such as microwave leakage) are present, q and q' should be equal. The unknown phase difference ϕ can then be estimated by equating (86) and (87) to find

$$\phi = \frac{u - u'}{E + E'} \pm \delta\phi. \quad (88)$$

The uncertainty

$$\delta\phi = \sqrt{\sigma^2 + \sigma'^2}/(E + E') \quad (89)$$

is typically about 0.6 % of ϕ , or $1.5 \mu\text{rad}$ compared with our 0.25 mrad value of ϕ . Over six years of measurements ϕ has been stable to about 1.5 % of its value, with no long-term trend like that observed with NBS-6 [41]. Figure 9 shows the history of the

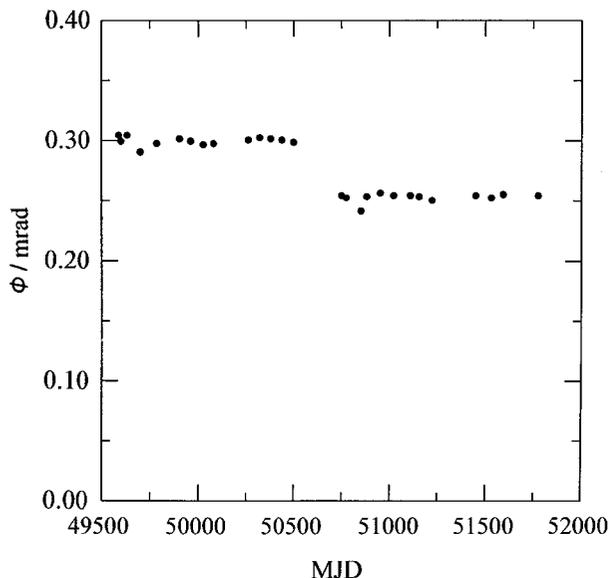


Figure 9. History of end-to-end cavity phase difference.

end-to-end cavity phase difference. The abrupt change in the summer of 1997 (MJD 50659) occurred when an internal microwave leak was closed.

The phase difference ϕ can be eliminated by forming a weighted average \bar{u} of the measurements in the two beam directions:

$$\begin{aligned} \bar{u} &\equiv \frac{E'u + Eu'}{E + E'} \\ &= q \pm \frac{\sqrt{E'^2\sigma^2 + E^2\sigma'^2}}{E + E'}. \end{aligned} \quad (90)$$

The frequency measurement uncertainties in \bar{u} are reduced by the combination of the two runs, just as if they had been made under identical conditions. If E and E' are comparable, then the coefficients of σ and σ' are approximately 0.5. If σ and σ' are also comparable, the uncertainty in \bar{u} becomes $\sigma/\sqrt{2}$.

Our source of leverage for the end-to-end phase bias comes from the fact that the accuracy of the elimination of ϕ does not depend on the frequency uncertainty, but is limited by the accuracy of the coefficients E and E' . If E and E' are in error by δE and $\delta E'$, then \bar{u} will be in error by

$$\delta\bar{u} = \left(\frac{\delta E'}{E'} - \frac{\delta E}{E} \right) \frac{EE'}{(E + E')} \phi, \quad (91)$$

which has an order of magnitude of $\phi\delta E$. From the accuracy of the velocity distributions and power scale we expect fractional uncertainties $\delta E/E$ of 0.1 % to 0.2 %. Uncertainties of this magnitude would translate into an uncertainty in \bar{u} of about $0.7 \times 10^{-15} \nu_{\text{hfs}}$, smaller than $E\delta\phi$ or the measurement uncertainty in (90). More probable is that errors in power level or velocity distribution will affect E and E' in the same way. That is, their uncertainties will be highly correlated. Then $\delta E/E$ and $\delta E'/E'$ will partially cancel in (91) and $\delta\bar{u}$ will be even smaller. However, we still assign a Type B uncertainty of $0.7 \times 10^{-15} \nu_{\text{hfs}}$ to the end-to-end phase bias.

Note that when an evaluation is performed with frequency measurements made in only one beam direction and a previous value for ϕ is used, we must assign the larger uncertainty $E\delta\phi \approx 4 \times 10^{-15}$.

When more than one measurement in a given beam direction is available, the corresponding frequencies and E values can first be averaged before applying (90). In the past we have used a trial value of ϕ and adjusted it to minimize the scatter in the resulting values of q and q' . This procedure gives a result differing from (90) by at most $10^{-15} \nu_{\text{hfs}}$.

4.5 Distributed-cavity phase bias

Because the microwave cavity has small losses in its walls, the phase of the microwave field varies slightly across the beam aperture. If the atomic beam illuminates

the aperture uniformly, then the average phase over the aperture is the relevant phase for computing the end-to-end phase difference. But if the atomic beams in the two directions illuminate the aperture differently, an additional bias occurs because the average phase seen by one beam differs from that seen by the opposite beam. This bias is called distributed-cavity phase.

By the design of the ring-shaped De Marchi cavity ends [17] the phase variation across our 3 mm circular aperture is estimated to be only a few microrads. If the field mode is not centred on the aperture, a left-right asymmetry is present. No vertical asymmetry is expected because the mode is fixed in that direction by the cavity walls. Let ϕ_L and ϕ_R be the average phases over the left and right halves of the aperture, respectively. We have measured the difference between ϕ_L and ϕ_R by blocking one half of the aperture with a movable mask, measuring the resonance frequency, and repeating with the other half blocked. The difference between the two measured frequencies is just the coefficient E times $\phi_L - \phi_R$. We have found a phase difference of 7.6 μrad for each cavity end with a 35 % uncertainty due to the uncertainty in the frequency measurements. Separate determinations of E for the left and right half apertures showed a change of only 0.7 % compared with the 3 % change in $E\phi$.

We have also measured the uniformity of illumination of the apertures by measuring the relative signals when one or the other mask was in place. Let f_L and f_R be the fractions of the atomic beam flux passing through the left and right halves of the aperture when illuminated by the far oven. Let f'_L and f'_R be the corresponding fractions when the aperture is illuminated by the nearer oven, but with the atoms still required to pass through the more distant full aperture. The average phase over the full aperture is then $\phi = f_L\phi_L + f_R\phi_R$ for the far oven, and $\phi' = f'_L\phi_L + f'_R\phi_R$ for the near oven. The distributed-cavity phase bias arises from the phase difference $\phi - \phi'$. It vanishes if $f_L = f_R$ and $f'_L = f'_R$ (equal illumination), or if $f_L = f'_L$ and $f_R = f'_R$ (perfect retrace of beams), or if $\phi_L = \phi_R$ (phase symmetry). The effect occurs independently for each cavity end. Measurements of the transmitted beam fractions with the ovens installed in September 1999 showed that the largest asymmetry was $f'_L = 0.39$ and $f'_R = 0.61$ for the east cavity end.

Combining the beam fraction measurements with the phase difference measurements gave $\phi - \phi' = -0.8 \mu\text{rad}$ for the east end of the cavity and $\phi - \phi' = -0.3 \mu\text{rad}$ for the west end. Note that the phase difference measurements were made with $f_L = 1, f_R = 0$ and the reverse, in contrast to the observed asymmetry. This is the source of leverage for this bias.

When frequency measurements in two beam directions are combined by (90) to eliminate the end-to-end cavity phase difference, we are left with a

distributed-cavity phase bias of the form

$$\frac{EE'}{E + E'}(\phi - \phi') \quad (92)$$

for each cavity end. Combining our computed values of E and E' with the measured phase differences leads to a net fractional frequency bias of -1.3×10^{-15} , with an uncertainty of 0.4×10^{-15} . The uncertainty is dominated by the Type A uncertainty in the frequency measurements.

These measurements should be repeated any time the ovens are changed or the beam alignment is otherwise affected. Our present results are not applicable to evaluations before September 1999, but the order of magnitude of the bias should have been similar.

4.6 Gravitational red shift

This bias arises from the theory of general relativity. An approximate solution of the theory's equations valid in the vicinity of the Earth equates the relative shift to the local gravitational potential. Since the TAI second is defined in terms of the rate of a clock on the Earth's geoid, the bias is due only to the difference in gravitational potential between the geoid and the location of NIST-7. This difference is adequately given by the model

$$\delta\nu_G = (gh/c^2)\nu_{\text{hfs}}, \quad (93)$$

where g is the local acceleration due to gravity and h is the height of NIST-7 above the geoid. The bias is computed to be 179.9×10^{-15} with an uncertainty of 0.1×10^{-15} [42]. It is relatively large because of the altitude of the Boulder Laboratories where NIST-7 rests. This bias should be constant over the operational life of the standard (barring a major earthquake). The bias is positive because the gravitational potential of the Earth is less negative at Boulder than on the geoid.

4.7 Black-body shift

This bias arises from non-resonant virtual excitation of the hyperfine levels to electronic levels (ac Stark shift) by the black-body background radiation bathing the atoms from the interior walls of the beam tube. Our model is the formula calculated by Itano et al. [43], adjusted for the improved measurement of the Stark effect made by Simon et al. [44],

$$\delta\nu_B = -1.711 \times 10^{-14} W^4 [1 + 0.014 W^2] \nu_{\text{hfs}}, \quad (94)$$

where W is the ratio of the absolute background temperature to 300 K. The accuracy of (94) is about 0.2 %, limited by the input value for the dc atomic polarizability.

For NIST-7 the internal temperature, 39 °C, gives a fractional frequency bias of -20.4×10^{-15} . The

uncertainty in internal temperature of 1 °C gives an uncertainty of 0.3×10^{-15} in this bias. The internal temperature is controlled so that $\delta\nu_B$ remains stable through many evaluations.

5. Frequency biases for which no corrections are made

The following effects are also potential sources of frequency bias. For each we estimate the associated bias and show that it is too small to require a correction. We also estimate an uncertainty for each bias to include in our overall Type B uncertainty.

5.1 Magnetic field inhomogeneity

There are two mechanisms resulting in three distinct biases introduced by magnetic field inhomogeneity. We first present our notation for describing inhomogeneity. Then we deduce formulae for each bias followed by quantitative estimates.

We write the strength of the longitudinal magnetic field

$$H(z) = H_0[1 - \epsilon f(z)], \quad (95)$$

where H_0 is the constant nominal value of the field, ϵ is a small constant describing the order of magnitude of the inhomogeneity, and the function f of the longitudinal position z is approximately unity or less. In this section only we use angle brackets to denote an average over the z values encompassing the drift region. The average drift region field is then

$$\langle H(z) \rangle = H_0[1 - \epsilon \langle f(z) \rangle], \quad (96)$$

as all the variation is contained in $f(z)$.

In the following we write ν_z for the Zeeman frequency defined by (5) with H_0 instead of H and $\delta\nu_{QZ}$ for the corresponding quadratic Zeeman effect defined by (72). That is, ν_z and $\delta\nu_{QZ}$ are the values when there is no inhomogeneity ($\epsilon = 0$).

5.1.1 Mean-square versus square-of-the-mean fields

The field-dependent transition frequency ν_1 defined by (6) is dominated by the linear Zeeman effect. The latter's measured value becomes

$$\nu_z(\text{meas}) = \nu_z[1 - \epsilon \langle f(z) \rangle]. \quad (97)$$

If this measured value is used to calculate the quadratic Zeeman effect of the clock transition we obtain

$$\begin{aligned} \delta\nu_{QZ}(\text{calc}) &= 8\nu_z(\text{meas})^2 / \nu_{\text{hfs}} \\ &= \delta\nu_{QZ}[1 - 2\epsilon \langle f(z) \rangle + \epsilon^2 \langle f(z) \rangle^2]. \end{aligned} \quad (98)$$

The quadratic Zeeman bias of the clock transition is proportional to the average of the square of the magnetic

field $\langle H(z)^2 \rangle$ seen by the atoms. It can be represented by

$$\delta\nu_{\text{QZ}}(\text{atom}) = \delta\nu_{\text{QZ}}[1 - 2\epsilon\langle f(z) \rangle + \epsilon^2\langle f(z)^2 \rangle]. \quad (99)$$

The difference between (98) and (99) is the bias

$$\delta\nu_{\text{QZ}}(\text{square}) = \delta\nu_{\text{QZ}}(\text{calc}) - \delta\nu_{\text{QZ}}(\text{atom}) \quad (100)$$

$$= \epsilon^2\delta\nu_{\text{QZ}}[\langle f(z)^2 \rangle - \langle f(z) \rangle^2]. \quad (101)$$

Because this bias depends on the square of the inhomogeneity ϵ , it becomes very small for standards with good field homogeneity. For NIST-7 ϵ is about **0.0005**. The square bracket in (101) is at most 0.5, but about **0.1** for the measured inhomogeneity [15]. Hence for NIST-7 with $\delta\nu_{\text{QZ}} = 1.4$ Hz, this bias is expected to be about 4×10^{-8} Hz, or $4 \times 10^{-18}\nu_{\text{hfs}}$, but certainly not more than $2 \times 10^{-17}\nu_{\text{hfs}}$. This bias is independent of microwave power and modulation.

5.1.2 Pedestal-induced bias of clock transition

The second mechanism creating a frequency bias arises from field inhomogeneity where the magnetic field in the excitation regions differs from the mean field in the drift region. This shifts the Rabi pedestal away from the centre of the Ramsey fringe. From (39) the Rabi pedestal is a symmetric sum of the excitation probabilities in the two excitation regions. The pedestal of the clock transition then has the quadratic Zeeman effect

$$\delta\nu_{\text{QZ}}(\text{ped}) = \delta\nu_{\text{QZ}}[(1 - \epsilon f_1)^2 + (1 - \epsilon f_2)^2]/2, \quad (102)$$

where f_1 and f_2 are the average values of $f(z)$ in the excitation regions. From (99) and (102) the pedestal offset of the clock transition is

$$\begin{aligned} \delta\nu_0(\text{off}) &= \delta\nu_{\text{QZ}}(\text{ped}) - \delta\nu_{\text{QZ}}(\text{atom}) \\ &= \epsilon\delta\nu_{\text{QZ}}(2\langle f(z) \rangle - f_1 - f_2), \end{aligned} \quad (103)$$

where ϵ^2 terms are neglected.

To estimate this shift we measure the pedestal offset of the field-dependent $m = 1$ transition. The linear Zeeman effect for the pedestal is $\nu_{\text{Z}}(1 - \epsilon f_1/2 - \epsilon f_2/2)$, compared with the Ramsey fringe effect given in (97). The difference is the pedestal offset

$$\delta\nu_1(\text{off}) = \epsilon\nu_{\text{Z}}(\langle f(z) \rangle - f_1/2 - f_2/2). \quad (104)$$

From (103), (104), and the definition of $\delta\nu_{\text{QZ}}$ (72), we can write the pedestal offset for the clock transition in terms of the pedestal offset of the field-dependent transition

$$\delta\nu_0(\text{off}) = 16(\nu_{\text{Z}}/\nu_{\text{hfs}})\delta\nu_1(\text{off}). \quad (105)$$

Measurements with NIST-7 show that $\delta\nu_1(\text{off})$ is about 20 Hz when the shields have been carefully degaussed,

and ν_{Z} is about 40 kHz. Combining these in (105) we find $\delta\nu_0(\text{off})$ to be about 1.4 mHz, which is too small to easily measure.

From the theory given in Section 3.3.3 a pedestal offset induces a bias of the Ramsey fringe given by (51). In the present case it becomes

$$\delta\nu_0(\text{inhomo}) = (\ell/L)\Lambda_0\delta\nu_0(\text{off}). \quad (106)$$

This bias should be added to $\delta\nu_{\text{QZ}}(\text{atom})$ in (99). The coefficient Λ_0 is of order unity and depends on microwave power, modulation amplitude, and velocity distribution, as shown in (52). For NIST-7, $\ell/L = 0.015$ and $\Lambda_0 = 1.3$ at 2.5 dB below optimum power. Combining these numbers with the estimate for $\delta\nu_0(\text{off})$, the bias on the Ramsey fringe is $3 \times 10^{-15}\nu_{\text{hfs}}$.

In the preceding discussion we ignored the pedestal offset caused by cavity pulling. For the clock transition pedestal this offset dominates, but for the $m = 1$ transition it is smaller than the inhomogeneity bias. It can be eliminated approximately by also measuring the pedestal offset for the $m = -1$ transition, then halving the difference between the two pedestal offsets. The cavity-pulling pedestal offset is nearly the same for the two transitions, and so drops out of the difference [40].

5.1.3 Pedestal-induced bias of field-dependent transition

The pedestal shift $\delta\nu_1(\text{off})$ also induces a bias in the Ramsey fringe for the $m = 1$ transition

$$\delta\nu_1(\text{inhomo}) = (\ell/L)\Lambda_1\delta\nu_1(\text{off}). \quad (107)$$

For $\delta\nu_1(\text{off}) = 20$ Hz, $\delta\nu_1(\text{inhomo})$ is about 0.4 Hz. If no correction is made for this bias, it becomes part of our experimentally measured Zeeman frequency. That is, $\delta\nu_1(\text{inhomo})$ must be added to $\nu_{\text{Z}}(\text{meas})$ in (97). The quadratic Zeeman effect calculated from $\nu_{\text{Z}}(\text{meas})$ then has a term $16\nu_{\text{Z}}\delta\nu_1(\text{inhomo})/\nu_{\text{hfs}}$ added to (98). This term is approximately $3 \times 10^{-15}\nu_{\text{hfs}}$ in our example.

5.1.4 Summary of inhomogeneity biases

Comparing the expressions for $\delta\nu_1(\text{inhomo})$ and $\delta\nu_0(\text{inhomo})$ we find the sum of the inhomogeneity biases

$$\begin{aligned} \delta\nu_{\text{QZ}}(\text{inhomo}) &= \delta\nu_{\text{QZ}}(\text{square}) + \\ &16(\nu_{\text{Z}}/\nu_{\text{hfs}})(\ell/L)(\Lambda_1 - \Lambda_0)\delta\nu_1(\text{off}). \end{aligned} \quad (108)$$

The last two terms have opposite signs. The two values of Λ differ only because the matrix element for the $m = 1$ transition, and hence the associated b value, is 3 % smaller than for $m = 0$. Actual computations at a power 2.5 dB below optimum show that the Λ values differ by only 1 %. Thus, the two pedestal offset biases nearly cancel, leaving a net effect of only about $3 \times 10^{-17}\nu_{\text{hfs}}$. If desired, it could be computed with an uncertainty of 10 %.

Because bias 1 is so small and biases 2 and 3 so nearly cancel each other, no correction for magnetic field inhomogeneity is included in the reported evaluations of NIST-7. We assign an uncertainty equal to the net uncorrected bias of $3 \times 10^{-17} \nu_{\text{hfs}}$.

5.2 Rabi pulling

Rabi pulling refers to the bias caused by the overlap of a transition lineshape by the tails of the Rabi pedestals of the adjacent transitions in the Zeeman spectrum. From the theory in Section 3.3.2, the bias is given by

$$\delta\nu_R = 2Y\nu_{\text{mod}}/D, \quad (109)$$

where D is twice the slope of the lineshape being measured at the modulation amplitude ν_{mod} , and Y is the slope of the tails of the overlapping transitions. If the Zeeman spectrum is symmetric, Rabi pulling of the clock transition vanishes. Because the wings of the Rabi pedestal decrease as detuning to the inverse fourth power for half-sine-wave excitation (24c), and little asymmetry is observed in the Zeeman spectrum, Rabi pulling is very small for NIST-7.

Like cavity pulling, (109) holds for pulling of either the Ramsey fringe or the Rabi pedestal. As the pedestal slope is much smaller and the modulation amplitude is much larger, Rabi pulling is very much larger for the Rabi pedestal than for the Ramsey fringe. By measuring offsets of pedestals from their Ramsey fringes, we obtain leverage for observing Rabi pulling. We obtain further leverage by making observations at lower C-fields to enhance the overlap, and at non-zero m values, where the neighbouring transitions are not symmetric.

The offset of the Rabi pedestal from its Ramsey fringe is dominated by three biases: cavity pulling, field inhomogeneity, and Rabi pulling. These effects can be separated by their m dependence [40]. Cavity pulling is approximately independent of m , but has a small even dependence due to the change in effective b value with the transition matrix elements. The inhomogeneity contribution is linear in m from the linear Zeeman effect. For Rabi pulling we note that the pedestal tails, given by (24c), are proportional to b^2 as their excitation is weak. From the transition matrix elements b^2 is proportional to $1 - m^2/16$. From this dependence we deduce that the net slope Y of overlapping transitions on both sides is proportional to m . However, the pedestal slope D_m in the denominator decreases with m , strongly for $m = \pm 3$. The offset of the m -transition pedestal can then be written

$$\delta\nu_m(\text{off}) = \delta\nu_m(\text{cav}) + m\epsilon\nu_Z((f(z)) - f_1/2 - f_2/2) + mB/D_m, \quad (110)$$

where the inhomogeneity term comes from (104) and B is a constant independent of m . If we average the offsets for $+m$ and $-m$, we obtain the cavity pulling

contribution. If we take the difference between the offsets for $+m$ and $-m$ and then divide by $+m$, we are left with a constant inhomogeneity contribution plus a Rabi pulling contribution with the m -dependence of the inverse pedestal slope. Using calculated (or measured) values of the slope we have separated the inhomogeneity and Rabi pulling contributions to (110) to find B . As an extreme example, the measured pedestal offset of the $m = +3$ transition at a C-field for which $\nu_Z = 19.7$ kHz was 302 Hz. We ascribe 3 Hz to cavity pulling, 108 Hz to field inhomogeneity, and 191 Hz to Rabi pulling. At our normal C-field, $\nu_Z = 39.5$ kHz, these values are 2 Hz cavity pulling, 53 Hz field inhomogeneity, and 2 Hz Rabi pulling.

Once B is determined by such analysis of pedestal offsets, we can estimate Rabi pulling of the Ramsey fringe by substituting the appropriate slope and modulation amplitude in (109). For our extreme example we estimate a fringe bias of 2 mHz, less than the noise on the $m = 3$ transition measurement.

For the clock transition to be biased by Rabi pulling we must have an asymmetry between the $m = +1$ and $m = -1$ transitions. We have observed asymmetries in peak signal heights up to 1 %. Assuming 1 % asymmetry as a worst case, we estimate fractional Rabi pulling of the clock transition pedestal to be 1.5×10^{-10} at $\nu_Z = 19.7$ kHz. The corresponding fractional pulling of the Ramsey fringe is 2×10^{-15} . These small values assure us that we could take frequency data at this lower field, or any intermediate field, if we wished. At our normal C-field, 1 % asymmetry gives 1.5×10^{-13} fractional pulling of the clock transition pedestal, but only 4×10^{-17} fractional pulling of the clock transition fringe. We assign a 50 % uncertainty to this estimate based primarily on uncertainty in the asymmetry assumed.

5.3 Ramsey pulling

Ramsey pulling arises from weakly excited $\Delta m = \pm 1$ transitions to or from an $m = 0$ sublevel occurring in one excitation region, while the clock transition occurs in the other. It vanishes if the $\Delta m = \pm 1$ transitions are not excited or if the Zeeman spectrum is symmetric. A detailed calculation of Ramsey pulling is difficult [45] and has not been carried out for the microwave field geometry used in NIST-7. But we use this theory to make rough estimates.

When $L \gg \ell$ the dominant term in Ramsey pulling is the third one in equation (60) of [45]. It arises from weak, non-resonant transitions from the $F = 3$, $m = \pm 1$ sublevels to $F = 4$, $m = 0$, followed by the clock transition. In our notation it becomes

$$\lambda_F = -(3b^2/2\omega_Z^2) \tan^2 \theta \times \left\langle \sin 2b\tau \sin b\tau \sin \left(\frac{1}{2} \omega_Z \tau \right) \sin \omega_{\text{mod}} T \right\rangle \times (N_1 - N_{-1})/N_0 D_0. \quad (111)$$

The factor $3/2$ comes from the matrix element for the $F = 3, m = \pm 1$ to $F = 4, m = 0$ transition, θ is the angle between the microwave field vector and the C-field, and the N_m are the initial populations of the $F = 3$ sublevels. We have assumed that θ is small and $b \ll \omega_z$. These assumptions are valid for NIST-7.

Another contribution arises from the clock transition followed by transitions to the $F = 4, m = \pm 1$ sublevels. Its formula is identical to (111) except that $3/2$ is replaced by $5/2$ and the relative populations are replaced by relative detection efficiencies.

As a function of Zeeman frequency, (111) represents a damped oscillation. The period is twice that predicted for Rabi pulling [46]. For a single velocity the damping is slow, due only to the coefficient. For a broad velocity distribution the damping is rapid due to the averaging of $\sin(\omega_z \tau / 2)$. As a function of microwave power, (111) increases rapidly with a slow oscillation. Dependence on the velocity distribution is significant.

For a worst-case estimate we consider ν_z about 20 kHz, optimum excitation power, $\theta = 0.16$ and $(N_1 - N_{-1})/N_0 = 0.02$. Ramsey pulling would then be $2.4 \times 10^{-15} \nu_{\text{hfs}}$ for NIST-7 with single velocity atoms. However, this estimate decreases to $0.6 \times 10^{-15} \nu_{\text{hfs}}$ when the velocity average is included. The C-field at which NIST-7 has usually operated ($\nu_z = 39.5$ kHz), happens to be close to a zero-crossing for Ramsey pulling. At this C-field we estimate Ramsey pulling to be only $2 \times 10^{-18} \nu_{\text{hfs}}$. In typical operation Ramsey pulling may be five times smaller still, due to lower microwave power and smaller asymmetry. The small size of the cavity beam hole (small θ), the symmetry of the spectrum, and the broad velocity distribution all combine to make Ramsey pulling much smaller for NIST-7 than for more conventional standard designs.

5.4 Bloch-Siegert shift

This bias arises from the “anti-rotating” component of the applied microwave field inducing the transitions [47, 48, 28]. It is unavoidable except for magnetic resonance, where circularly polarized excitation can physically correspond to a “rotating” wave. It has the same form as biases due to non-resonant excitation of other atomic states by the microwave field, or biases due to non-resonant sidebands on the microwave signal. In our notation the Bloch-Siegert shift is given simply by b^2/ω_0 and is positive.

For frequency standards using Ramsey excitation the Bloch-Siegert shift was computed long ago for single-velocity atoms, constant excitation amplitude, and no modulation [27]. As this bias occurs only in the excitation regions, it can be treated by the analysis for the inhomogeneity shift in Section 3.3.3. For constant excitation we have from (51) and (52)

$$\lambda_{\text{BS}} = (b^2/\omega_0)(\ell/L) \times \langle h(\tau) T \sin^2 2b\tau \sin \omega_{\text{mod}}(T + h\tau) \rangle / D. \quad (112)$$

For half-sine-wave excitation the time-dependence must be included. Within the excitation region the Bloch-Siegert shift becomes

$$\delta\omega_{\text{BS}} = (b_0^2/\omega_0)(\pi^2/4) \sin^2(\pi t/\tau). \quad (113)$$

When this time-dependence is integrated in (28a) we obtain in place of (112)

$$\lambda_{\text{BS}} = (b^2/\omega_0)(\ell/L) \langle T [J_0(a) + J_2(a)] \text{sec } a \sin^2 2a \sin \omega_{\text{mod}}(T + h\tau) \rangle / D, \quad (114)$$

where $a = b_0\tau$. That is, $h(\tau)$ is replaced by $[J_0(a) + J_2(a)] \text{sec } a$. For typical power levels (114) is about 25 % greater than (112).

For NIST-7 at -2.5 dB below optimum power the Bloch-Siegert bias is $3 \times 10^{-16} \nu_{\text{hfs}}$, increasing to $8 \times 10^{-16} \nu_{\text{hfs}}$ at 3 dB above optimum power. The uncertainty would be about 1 % of the bias arising from the uncertainty in the velocity distribution and associated b value.

As caesium standards have evolved over the years the Bloch-Siegert bias has always remained negligible. Designs increasing the line Q by increasing L reduce the ℓ/L factor. Designs using slower atoms operate with a smaller b as τ is larger. Both factors decrease λ_{BS} .

5.5 Fluorescent light shift

There are two sources of light-shift bias arising from near-resonant excitation of optical transitions in the caesium atoms. One is caused by fluorescent light emitted along the atomic beam direction by optically excited atoms in the state selection or detection process. The other comes from laser light mechanically scattered into the drift region. Our model for the first is based on the theoretical work of Shirley [49] which, for the accuracy required here, has been verified on two other thermal beam machines [50, 51]. The input data are the geometry of the standard and the operating conditions. Under our normal operating conditions the predicted shift is 3×10^{-17} , with an uncertainty of no more than a factor of two. There is no change we can make to our operating conditions (pumping zone, pumping transition, light intensity or polarization or even two laser pumping) that will make the predicted shift measurable. One relevant experiment was conducted during the search for collisional effects (Section 5.7). We operated the standard with atomic beams and their associated optical pumping simultaneously present in both beam directions. This introduced a second major source of scattered light. As expected from the model, we did not see a shift in our frequency comparison greater than the Type A uncertainty of a few parts in 10^{15} .

Mechanically scattered light is essentially impossible to model, as one cannot quantify the location and magnitude of the scattering sites. However, it is easy

to vary the laser beam's intensity over a large range (a factor of 10) without affecting the clock operation. We have performed these tests and see no effect at the level of a few parts in 10^{15} . Hence, with our normal beam intensities we assign a relative uncertainty of 0.5×10^{-15} .

A recently introduced technique to measure the effects of light shift takes the form of a parametric test [50]. One looks for a change in the frequency of the standard as the microwave power and modulation depth are varied. Following our philosophy of the evaluation process in Section 1.1.3, we hesitate to use such a test for quantitative analysis of any bias because it requires a complete, quantitative knowledge of every other bias that depends on the altered parameters. For example, we have a small amount of microwave leakage (Section 5.9) within the beam tube whose effects are cancelled in the beam reversal process but whose power dependence is not adequately modelled. For this reason we have not used the test.

5.6 Majorana transitions

Majorana transitions are transitions between the Zeeman sublevels of either of the hyperfine levels. They can be caused by inhomogeneous magnetic fields that a moving atom sees as rotating around the dominant C-field. They can easily occur in standards with magnetic state selection or when atoms pass through holes in magnetic shields. Majorana transitions cause shifts only if they create coherences among the Zeeman sublevels in a manner asymmetric in m . These coherences must then interact with a microwave excitation field that has a component perpendicular to the C-field. The design of NIST-7 with the entire state-selection, excitation and detection process performed in a single region of highly uniform magnetic field, combined with the high symmetry of sublevel populations achieved with optical pumping, make Majorana transitions negligible.

A rough estimate of the possible bias from Majorana transitions was made based on the theory of Bauch and Schröder [52]. The result was $1.3 \times 10^{-15} \nu_{\text{hfs}}$. This estimate included a 1 % asymmetry between the $m = 1$ and $m = -1$ populations and a factor 4 reduction for the velocity average of phases dependent on transit time. The bias must also be averaged over trajectories through the beam aperture. As beam flux measurements through the right and left halves of the beam aperture showed a 20 % asymmetry in one case, we have used a factor 5 reduction. We have made no restriction on the degree of coherence generated, so our estimate is still conservative.

Like Rabi and Ramsey pulling, a bias due to Majorana transitions would increase rapidly as the C-field is reduced due to the reduced detuning of the microwave π -transitions. The absence of any significant bias, other than those discussed in Section 4, when we operated at $\nu_Z = 19.7$ kHz, is further assurance that these biases are very small.

5.7 Collisional shifts

Collision-induced frequency biases have never been observed in thermal atomic beam standards. However, some qualitative arguments have been put forward [53, 29, 54] that suggest collisional effects might be significant at our level of accuracy. To our knowledge the frequency bias for caesium collisions has never been treated theoretically for the case of thermal atomic beams. We thus have no quantitative model for estimating the bias.

In Section 5.6.2(D) of [29] the authors present a model for a rough estimate of the bias due to scattering from background caesium atoms. We have carried out resonance fluorescence experiments to search for background caesium atoms within the beam tube. We operated one atomic beam optically pumped into the $F = 3$ state but with the microwave excitation turned off. The detection laser downstream was then moved around to look for atoms in the $F = 4$ state travelling in directions different to that of the beam atoms. The laser windows constrained us to look only in regions very close to the atomic beam. Within the limits imposed by the Lorentzian tails of the resonance from unpumped atoms in the beam, we could detect no background. This means that in the detection region the background density cannot be more than 1 % of that in the atomic beam, or 1×10^{-10} Pa. While we could not make measurements directly in the drift region of the beam tube, the geometry of the standard is such that the background caesium pressure in that region could not be orders of magnitude larger than in the detection region. Even using 1×10^{-8} Pa as an upper bound to the possible background pressure in the model of Vanier and Audoin [29], we arrive at a bias of only 7×10^{-16} for interactions with background caesium atoms.

An additional experiment was carried out operating NIST-7 with the opposing atomic beam active. This is possible because pumping the opposing beam to the $F = 3$ hyperfine level allows it to traverse the detection zone for the observed beam without fluorescing. The density in the beams is so low that we see no attenuation of the observed beam due to the opposing beam. Furthermore, no frequency bias was observed within the 4×10^{-15} Type A uncertainty of the measurements. If the collision bias obeys a semi-classical model (Chap. 3 in [29]), it is proportional to $n\bar{v}\sigma$, where n is the density, \bar{v} is the relative velocity, and σ is the cross-section. The opposing beam increases n by two and \bar{v} by about 5. So even if σ decreases by a factor of 2 at the higher relative velocities, the bias would have been increased by a factor of 5. However, the atoms experiencing the larger bias from opposing beam collisions may be scattered out of the beam and not detected. So the opposing beam experiment may only increase the background density without giving additional information.

Because of the speculative nature of the above arguments, we have assigned an uncertainty of 1×10^{-15} to the collisional bias.

5.8 Beam-flux variation

If the atom or laser beam flux, hence the signal level, is changing slowly and the frequency servo always measures the same side of the lineshape first, a frequency bias results. The bias is given simply by $(dS/dt)\delta t/KD$, where (dS/dt) is the time rate of change of the signal, δt is the time interval between measurements on the two sides of the line, and KD is twice the signal slope at the detuning of the modulation. A linear decrease of 10 % in signal over a 12-hour measurement causes a bias less than $10^{-15}\nu_{\text{hfs}}$. We have reduced this bias by controlling the temperature of the caesium oven so that the maximum observed change in the clock signal amplitude is less than 10 % over a 24-hour period. Furthermore, the frequency servo now alternates which side of the Ramsey fringe it measures first, eliminating the effect of slow linear changes in the signal. We assign an uncertainty of less than $10^{-16}\nu_{\text{hfs}}$ to this effect.

5.9 Microwave leakage

When caesium atoms interact with coherent excitation fields in regions outside the microwave cavity, the relative phases of the hyperfine amplitudes can be altered so that the centre of the Ramsey fringe is biased. Radiation from sources not related to the standard are not important because even a small difference in frequency causes the relative phase of the induced transition to quickly average to zero. Therefore, we need concern ourselves only with sources of radiation from the microwave synthesis chain, the delivery system external to the beam tube, and microwave structures within the beam tube. Radiation originating outside the beam tube travels through the laboratory over multiple, unstable pathways to ports where it couples into the beam tube. Because the phase of this radiation is uncontrolled, we find its presence totally unacceptable. With a sensitive, heterodyne detector we can detect and then reduce all such external sources of radiation to acceptable levels [55, 9]. The initially constructed microwave system induced shifts of order 10^{-13} . We have been able to reduce that leakage by 50 dB. As the leakage bias is roughly proportional to the square root of the leaked power, we believe that any bias from external microwave leakage is now less than 1×10^{-15} .

Within the beam tube, a low level of leakage from joints in the microwave plumbing is more tolerable. The design of NIST-7 allows beam reversal simply by changing the current to the oven heaters. The external location of the ovens keeps thermal gradients well away from the cavity. As no mechanical motion is involved, any leakage field remains stable in amplitude and phase. This argument is supported by the observed long-term stability of the end-to-end phase difference shown in Figure 9.

A leakage model shows that the leakage bias has the order of magnitude of the strength of the leakage field in units of Rabi frequency. Further, to produce a bias, the leakage field must be out of phase with the field inside the cavity and have an asymmetric spatial distribution along the atomic beam. Graphite pieces, placed in the beam tube to absorb caesium, also absorb microwaves, creating net travelling waves of leakage fields. Travelling waves have components out of phase with the field inside the cavity. Leakage from the centre of the microwave cavity or equally from joints at each end is symmetric and does not cause a bias. Leakage from one end does. But the bias caused by spatial asymmetry changes sign with reversal of the atomic beam direction.

Our model of the microwave leakage bias was partially verified experimentally by intentionally introducing leakage fields through an unused electrical feedthrough at one end of the vacuum chamber. We were able to introduce frequency biases of a few parts in 10^{11} . These were cancelled by beam reversal to better than 1 %. Thus, if the change in end-to-end phase difference seen in Figure 9 was due to a reduction of microwave leakage, our results prior to the change were in error by no more than 1×10^{-15} .

5.10 DC Stark effect

This bias arises if the atoms are subjected to a static (or low-frequency) electric field while their resonance is being measured. It is a quadratic effect due to the atomic polarizability. A recent measurement with a caesium fountain [44] gives the following relation

$$\delta\nu_S = -2.27 \times 10^{-10} E^2 \text{ Hz (V/m)}^{-2}. \quad (115)$$

Thus, a fractional frequency bias of 10^{-15} is realized with a field $E = 200 \text{ V/m}$. Since the internal structures of NIST-7 are separated by the order of 0.1 m in the drift region, a potential difference of 20 V would be required to generate such a field. Although the major parts of NIST-7 are insulated from each other to avoid current loops, they are all connected to a common ground at one place. This design creates a Faraday cage within which there should be no fields.

In spite of this careful design there are two possibilities for static fields to exist within the beam tube. One is charge accumulation on the insulating film used to isolate the cavity from its support structure. This material is an excellent insulator capable of supporting tremendous surface charge. However, all the insulating film is trapped between conducting metal surfaces with zero potential difference. The only exception to this is the few millimetres of material that may protrude around the edges. In the absence of any source of directed, high-energy electrons, we believe that there is no way for charge to accumulate on these pieces of exposed material. The random occurrence of charged particles cannot lead to a large charge build-up because

subsequent low-energy particles would be repelled by the smallest of charges on the insulating film. Furthermore, field lines from any such charge would terminate on the adjacent metallic surfaces and not extend across the beam line.

The second area of possible concern involves something called “patch” effects on the metal surfaces very near the atomic beam. The only such surfaces in NIST-7 are the holes in the microwave cavity through which the atomic beam passes. The base material is electroformed copper. However, it is covered with a rather thick layer of oxide. Investigations performed in relation to studies of the gravitational potential on charged particles [56] indicate that surface potentials of the order of **0.01 V** can be expected and these will vary on a spatial scale of order micrometres. This situation can generate field gradients of the order of 10^4 V/m but they extend only micrometres into the atomic beam. Furthermore, they exist over only about 1 % of the atomic beam length. Hence, fractional frequency shifts greater than 10^{-17} are not expected from this effect.

5.11 Biases related to electronics

To reach our target uncertainty of 5×10^{-15} we must determine the centre of the Ramsey fringe with an uncertainty less than one millionth of its width. We have identified two main sources of frequency biases that occur with the slow square-wave digital frequency servo: microwave sideband pulling and demodulator errors. We discuss each type of bias below.

5.11.1 Microwave spectrum

Discrete sidebands in the spectrum of the microwave radiation may produce errors in the measurement of the centre of the Ramsey resonance [57]. RF sideband pulling was studied using NIST-7 by intentionally modulating the microwave signal in both amplitude and phase to produce a well-defined spectrum [22]. To determine the magnitude of the rf sideband bias under operating conditions, it is sufficient to record the amplitude-modulated sideband spectrum alone [21].

Our measurements indicate that the rf sideband bias for NIST-7 is

$$\delta\nu_{\text{sideband}} = (0 \pm 0.1) \times 10^{-15}. \quad (116)$$

We have found that spurious sidebands are produced most often by improper grounding of the microwave signal path. Therefore, we discuss the grounding of the microwave synthesis chain electronics below.

Ground loops

We use the term *ground loop* to refer to the electric circuit formed by, as well as the current flowing in, a continuous, low-impedance path formed from a system of conductors. These conductors may include structural

metal parts of the chassis, equipment racks, and even the vacuum system, as well as cable shields. Oscillating magnetic fields, especially at the frequency of the power mains, can induce alternating currents within such a circuit. Electrical connections between subassemblies of the physics package are discussed in Section 2.1. Here we describe proper connection of the synthesizer components to the laboratory ground reference.

The microwave synthesis chain is divided into subsystems that are electrically isolated from one another at low frequencies while maintaining low insertion loss and low leakage at high frequencies. The **5 MHz** signal from the reference oscillator is transmitted to the primary standards laboratory from a high-isolation distribution amplifier [58] in another room. In addition, transformers are used to electrically isolate the **5 MHz** transmission line. The direct digital synthesizer (DDS) that provides slow square-wave modulation is optically isolated from the servo computer’s data bus. Its connection to the synthesis chain is isolated at low frequencies by a transformer. Following the synthesis chain, a computer-controlled attenuator is used to actively stabilize the microwave power level. As this circuit requires a very low servo bandwidth, it is isolated from the computer’s digital-to-analogue converter with a differential-input instrumentation amplifier. Isolating the beam tube from the power servo is, however, more difficult. Traditional devices that block direct current paths emit unacceptable amounts of radiation at **9 GHz**. We have developed a unique waveguide dc blocking device that presents an impedance of **2 k Ω** to low frequencies with an insertion loss of just a few tenths of a decibel at **9 GHz**.

5.11.2 Modulation-synchronous errors

The slow square-wave frequency servo determines the centre of the Ramsey lineshape using the resonance condition (46)

$$S_+(f_{\text{res}}) - S_-(f_{\text{res}}) = 0. \quad (117)$$

A bias occurs if the signal measured at one side of the lineshape is shifted by an amount δS with respect to the signal from the other side. With a change to the software, the digital servo can be used to measure biases of this type [23]. The integral controller described in Section 3.4 is replaced by a proportional controller. The servo loop is then interrupted in a way that both removes the frequency dependence of the error signal and emphasizes the bias of interest. The resulting measured frequency has the bias (62)

$$\delta\nu = -G\delta S \quad (118)$$

for a proportional servo using the ideal gain G . In order for this bias to be less than 1×10^{-15} , δS must not exceed $1 \mu\text{V}$, corresponding to 4×10^{-4} of one least significant bit (LSB) of the analogue-to-digital converter

(ADC) used to digitize the clock signal. This level of performance might seem unreasonable given a typical ADC differential non-linearity of ± 1 LSB. However, we have demonstrated the ability of our demodulator to measure error signals at this level [23].

Below, we present modifications of the servo loop that test for potential sources of demodulator bias. We stress that it is important to make as few changes as possible to the hardware and software during these bias measurements so that the results may be applied to the normal operation of the standard.

Demodulator bias

The digital servo performs many functions that are synchronous with the measurement of the Ramsey resonance. These tasks involve transmission of signals throughout the servo electronics. Therefore the fluorescence photodetector, preamplifier, and analogue-to-digital converter must be tested to ensure that they are not affected by the operation of the digital servo. These systems were tested by blocking both the atomic beam and the detection laser. The measured frequency bias from coupling of the modulation sequence to the photodetection and digitizing electronics is

$$\delta\nu_{\text{dcmoD}} = (0 \pm 1) \times 10^{-15}. \quad (119)$$

Laser amplitude-modulation bias

This bias occurs if the intensity of the detection laser changes synchronously with the frequency modulation by the frequency servo. To measure laser amplitude modulation the atomic beam is blocked. Laser light is then scattered into the detection region of the standard by an etched glass filter. Sufficient light is scattered to reproduce the nominal signal level. The measured fractional frequency bias from demodulator-correlated amplitude modulation of the laser is

$$\delta\nu_{\text{LaserAM}} = (0 \pm 1) \times 10^{-15}. \quad (120)$$

5.11.3 Switching transients

When the frequency servo software sends a modulation command to the DDS, the beam tube must be permitted to reach steady-state before the clock signal is measured. During the blanking interval τ_b the clock signal is discarded. Afterwards, the clock signal is recorded during the acquisition interval τ_a . If the frequency of the synthesis chain were to change instantaneously, the blanking interval would be selected so that the slowest atoms would have time to clear the Ramsey cavity. The subsequent measurement would involve only those atoms that had seen the same microwave frequency in each end of the Ramsey cavity.

However, the frequency does not change instantaneously. There are delays within the DDS due to processing of frequency commands. In addition there

are delays due to the settling time of the **10.7** MHz phase-locked loop [22] and the photodetector. These additional settling times extend the blanking interval to

$$\tau_b = \tau_r + \tau_\phi + \tau_p + \tau_d. \quad (121)$$

Here τ_r is the response time of the DDS, that is, the time from the transmission of a frequency command to the actual change in frequency at the DDS output. τ_ϕ is the phase-settling time of the **10.7** MHz phase-locked loop (Figure 3), τ_p is the transit time of the slowest atoms across the Ramsey cavity to the detection region, and τ_d is the photodetector's settling time. As the phase-locked loop steers the **10.7** MHz quartz oscillator to the new modulation point, the frequency sweeps over the peak of the Ramsey resonance. This produces a brief pulse of fluorescent light in the optical detection region. Time τ_d is chosen so that the resulting photocurrent transient settles prior to the acquisition interval. For the present configuration of the digital servo, $\tau_b = 60$ ms and $\tau_a = 1$ s.

Using the method of Section 5.11.2, we have measured the fractional frequency bias due to microwave switching transients. The most significant non-zero bias measured was

$$\delta\nu_{\text{ST}} = (-1.5 \pm 2) \times 10^{-15}. \quad (122)$$

This measurement was made on only one side of the lineshape, so may be partially or fully cancelled by the same effect on the other side of the lineshape.

6. Measurement statistics

In the preceding two sections we discussed at length our known biases and assigned their Type B uncertainties. In the context of (1) we have yet to examine the stability of the reference oscillator and the behaviour of our measurement noise. We also need to assign the Type A uncertainty for an evaluation. These topics are treated below.

6.1 Reference oscillator stability

Figure 10 shows typical Allan deviations for NIST-7, for a reference oscillator, and for the synthesis chain. The upper trace shows the stability of NIST-7 relative to a hydrogen maser reference oscillator. The middle trace shows the stability of the reference oscillator compared with an ensemble of five similar masers [59]. The lower trace shows the stability of the microwave synthesis chain. This latter trace was obtained by comparing two similar synthesizers driven by the same reference oscillator [21]. These data indicate that, for the duration of an accuracy evaluation or bias test (several days), the uncertainty of the frequency comparison is limited by the measurement noise of NIST-7. To ensure that this is true for any particular frequency comparison, we continuously compare the frequency of the reference oscillator with the maser ensemble.

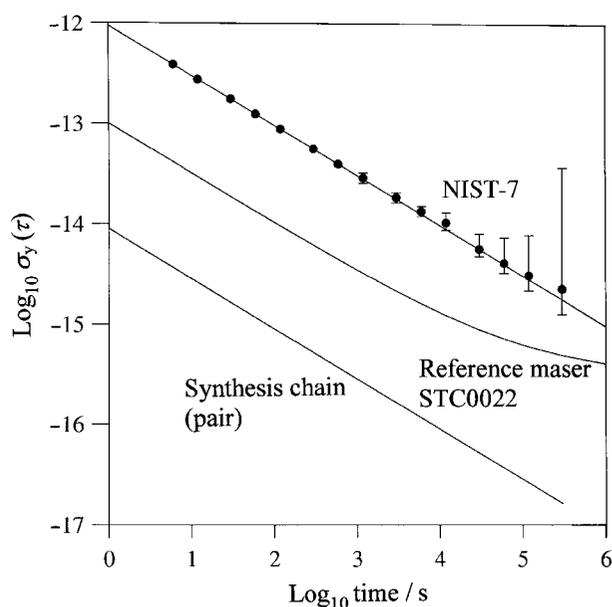


Figure 10. Comparison of Allan deviations for NIST-7, for a reference oscillator, and for the synthesis chain.

6.2 Noise character

The signals fed to the frequency servo contain noise. Sources of this noise include shot noise from the random arrival of atoms in the detection region; amplitude-modulated, frequency-modulated and pointing noise on the detection laser; noise in the photodetector, or in the analogue-to-digital converter; and the Dick effect [60]. Most of these sources can be controlled at a level below that of the atomic shot noise. Shot noise is expected to be white, i.e. to have no correlation with itself except at the same time [29]. To verify that our signal noise is white, we have measured the autocorrelation function of the error signal with the servo feedback to the synthesis chain disabled. The results, shown in Figure 11, are perfect correlation at zero time interval, and zero correlation for non-zero time intervals. These are the white-noise assumptions made in (61).

The theory in Section 3.4 predicts that the servo translates white error signal noise s_n into white frequency noise on the output time series f_n . This frequency noise does not reside in any electronic signal from the standard, but in the numbers f_n generated and stored by the servo computer. These numbers are analysed to find the mean frequency \bar{f} and its Type A uncertainty.

We routinely compute the Allan deviation (66) of the f_n and always observe a $\tau^{-1/2}$ slope, consistent with white frequency noise. This behaviour persists for measurement times of up to at least three days. A sample is shown in Figure 10. Although these results do not prove white frequency noise for all times, none of our tests has shown any departure from this characteristic.

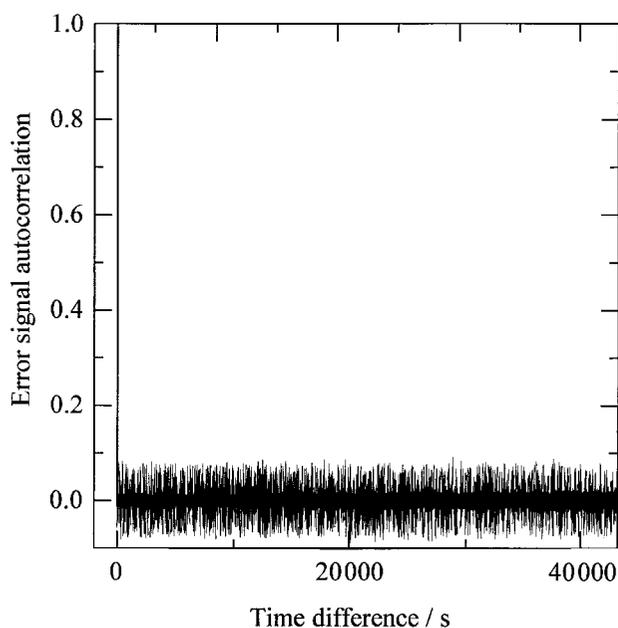


Figure 11. Autocorrelation of digital servo noise. Only the correlation at zero time difference is significant.

Dead time is introduced into our frequency measurements by the blanking time in signal observation and by the operation of the C-field and power servos. A special feature of white frequency noise is that the Allan variance is independent of any dead time in the measurements [38, 61]. No modification of the processing of the f_n is required to account for these dead times. Even the dead time required to reverse the atomic beam can be ignored.

6.3 Type A uncertainty

An estimate of the uncertainty in the mean frequency \bar{f} of a measurement comes from the variance of the mean

$$\sigma_{\bar{f}}^2 = \sigma_f^2 / N, \quad (123)$$

where σ_f^2 is the sample variance

$$\sigma_f^2 = \frac{1}{N-1} \sum_{n=1}^N (f_n - \bar{f})^2 \quad (124)$$

and N is the number of elements in the time series. This estimate corresponds to the Allan variance extended to the length of the time series. $\sigma_{\bar{f}}^2$ is unbiased only if the autocorrelation of s_n vanishes for $\tau > 0$ (see Figure 11) and if the digital servo employs the ideal servo gain G (64). In the past we have operated with a servo gain less than G . But now we measure G , with an uncertainty of no more than 5%, prior to each frequency measurement. The servo uses this gain for up to 12 hours, then measures it again. Thus, we may use (123) as an unbiased estimate of the Type A uncertainty for \bar{f} .

We can extend this analysis to multiple measurements that may not contain the same biases or Type A uncertainty. First we remove all known biases from each frequency measurement \bar{f}_k to obtain an unbiased frequency R_k . From a set of unbiased frequencies we can then compute a mean

$$\bar{R} = \frac{1}{N_e} \sum_{k=1}^{N_e} w_k R_k, \tag{125}$$

where N_e is the number of frequency measurements in the set and the optional w_k are relative weights for the R_k , assigned according to their respective uncertainties. Our estimate for the Type A uncertainty for \bar{R} is given by

$$\sigma_{\bar{R}}^2 = \frac{1}{N_e} \sum_{k=1}^{N_e} w_k^2 \sigma_{R_k}^2, \tag{126}$$

where

$$\sigma_{R_k}^2 = \sigma_{\bar{f}_k}^2. \tag{127}$$

$\sigma_{\bar{R}}$ is the Type A uncertainty we then assign to the mean frequency \bar{R} .

As a test of the variance $\sigma_{\bar{R}}^2$ we collected sets of up to fifty unbiased frequencies R_k from 12-hour measurements that had similar variances. These sets included measurements made over several months. From each group of measurements made only a few days apart (part of the same evaluation) we first removed the mean \bar{R} of those measurements to leave nominally random residuals. We thereby eliminated long-term drift of the reference oscillator from the residuals in a set. The variance of these residuals was then within 10 % of the variance $\sigma_{\bar{R}}^2$.

7. Evaluation process

An evaluation of NIST-7 is a series of measurements that both compares the frequency of the standard with that of the reference oscillator (hydrogen maser) and determines the biases. A complete evaluation determines all known biases, and includes tests for unexpected biases. This is a never-ending task. In practice we conduct routine evaluations that determine only the major biases that may change from one evaluation to the next. Between routine evaluations we make measurements to reduce the uncertainty of known biases and search for previously unknown biases.

7.1 Routine evaluations

A routine evaluation is carried out over a period of several days and repeated every sixty days. It includes several frequency measurements plus the auxiliary measurements for determination of the biases described below. The results of a routine evaluation are reported to the BIPM and published in the Time Section's *Circular T* [62].

7.1.1 Frequency measurements

A frequency measurement is made in a continuous run extending for many hours (typically 12) to reduce the Type A uncertainty. The result is a time series f_n of several thousand numbers, each representing a single value of the frequency difference $\Delta\nu$ defined by (1).

When a measurement has been completed, the mean \bar{f} and standard deviation σ_f of all the numbers is computed. There may be a few f_n quite far from the mean. This can be caused by, for example, a cosmic ray passing through the detector photodiode. If any f_n deviate by more than 5 standard deviations from the mean, they are discarded and the mean and standard deviation are recomputed. (We rarely see more than 3 discarded numbers in a set of about 14 000.) The Allan deviation of the data is also calculated and plotted as a routine check of the noise type.

Throughout the frequency measurement process the C-field is determined by measuring the frequency ν_1 of the $m = 1$ field-dependent transition every 30 s. The C-field servo holds ν_1 to a constant value that is recorded in the data file along with the servo corrections and the f_n . The frequency ν_1 is used to calculate the quadratic Zeeman bias as described in Section 4.1. Most evaluations have been made at a C-field of 5.64 μ T, corresponding to $\nu_Z = 39.5$ kHz.

The microwave power is also periodically measured and then held constant by the digital power servo (Section 2.2.3). Our experience with the synthesis chain shows that one measurement every 5 min is sufficient to hold the uncertainty in the power below 0.02 dB.

We make most frequency measurements at a microwave power 2.5 dB below optimum power and a modulation amplitude of 27 Hz. Unlike standards using magnetic state selection, our broad velocity distribution allows the excitation probability (the $\sin^2 2a(\tau)$ factor in (46)) and the modulation amplitude (the $\sin \omega_{\text{mod}} T$ factor in (46)) to become the velocity-determining elements. At powers below optimum we preferentially excite slower atoms, which produce a narrower Ramsey fringe. As the microwave power is reduced the signal level decreases, but line narrowing dominates until several decibels below optimum power. From computer simulations using an experimental velocity distribution we find that the slope of the Ramsey fringe has a maximum at 2.2 dB below optimum power, while the stability is best at 4.2 dB below optimum power. These extremes are broad, so little is lost by operating at slightly different values. Another reason we operate below optimum power is that most power-dependent biases decrease with microwave power.

7.1.2 Ramsey lineshape measurements

The Ramsey fringe lineshape is recorded before and after each frequency measurement. The signal amplitude is measured at discrete excitation frequencies

starting at line centre and progressing in equal steps to a detuning where the side lobes have damped out. Only one side of the lineshape is recorded since the asymmetry is smaller than the noise in the measurements. The lineshape measurements are made at four different power levels: **+2.5 dB**; **0.0 dB**; **-2.5 dB**; and **-5.0 dB** relative to optimum power. From these measurements we extract a transit-time distribution and Rabi frequency by the Fourier-transform method [34]. The second-order Doppler correction $\delta\nu_D$, cavity pulling $\delta\nu_C$, and end-to-end phase shift coefficient \mathbf{E} are then calculated as described in Sections 4.2, 4.3 and 4.4.

7.1.3 Zeeman spectrum measurements

Frequency measurements of the lines in the Zeeman spectrum are made once with each beam reversal, at a microwave power **2.5 dB** below optimum. A brief (200 s) measurement of the centre frequency of the Ramsey fringe is made for each line. The centre frequency of the Rabi pedestal is then measured for 1000 s. This gives typical uncertainties in fractional frequency of **0.7×10^{-13}** for the Ramsey fringe and **1.4×10^{-11}** for the Rabi pedestal.

We have published a detailed analysis of pedestal offsets for the Zeeman lines in [40]. From the offset of the clock transition pedestal we obtain an estimate of the cavity detuning slope needed for determining cavity pulling (Section 4.3). From the average of the offsets for the $\mathbf{m} = 1$ and $\mathbf{m} = -1$ pedestals we obtain a measure $\delta\nu_1(\text{off})$ of the C-field inhomogeneity (Section 5.1). This is used not only to estimate inhomogeneity biases, but also to check how well the shields have been degaussed.

The signal at the peak of each line is also measured and recorded to give a quantitative measure of the asymmetry of the spectrum. Observed asymmetries between the $\mathbf{m} = 1$ and $\mathbf{m} = -1$ lines are less than 1 %.

7.1.4 Electronics tests

Prior to each routine evaluation we test for spurious amplitude and frequency modulation of the 9 GHz radiation. We also perform a search for external microwave leakage. In the past these tests were carried out only occasionally, but they have proved to be so sensitive to minor changes in the electronics that we now perform them before each evaluation.

7.1.5 Summary of evaluation results

For each frequency measurement \bar{f}_k of an evaluation we remove the sum of biases given by (71) to leave the unbiased frequency \mathbf{R}_k . The biases are not necessarily the same for all k . The C-field and velocity distribution often change a little between measurements, while the

end-to-end phase bias changes sign with beam reversal. From these unbiased frequencies we then compute a mean and variance by the relations (125) and (126). Recalling (1), we interpret $\bar{\mathbf{R}}$ as the frequency difference between the definition and the reference oscillator. We assign $\sigma_{\bar{\mathbf{R}}}$ as the Type A uncertainty associated with $\bar{\mathbf{R}}$. These results, along with a summary of the biases removed and the Type B uncertainty, are reported to the BIPM.

The uncertainty $\sigma_{\mathbf{R}_k}$ for a typical 12-hour measurement has varied from **5×10^{-15}** to **1×10^{-14}** , depending on the signal strength and the duty cycle of the servo. Our evaluations have contained from six to ten or more 12-hour measurements. Hence our reported Type A uncertainty has varied from **4×10^{-15}** to as little as **2×10^{-15}** .

7.2 Tests for additional biases

Few unambiguous methods exist for finding additional biases. Some special techniques, described in Section 5.11.2, will reveal specific electronic biases. But otherwise we are limited to parametric measurements: changing some operating condition and looking to see whether there is an unmodelled change in frequency.

Since over half of the known biases depend on microwave power, we make a few measurements at a higher power level, **2.5 dB** or **3.0 dB** above optimum power, as part of each routine evaluation. The second-order Doppler bias and the end-to-end phase bias each increase by about one part in **10^{13}** with this change in power. In the east-to-west beam direction these changes add, but in the west-to-east beam direction they nearly cancel each other. After removal of these two biases and cavity pulling the residual frequencies should be independent of microwave power if our models are correct. In the past we have found an unexplained power shift of up to 4 parts in **10^{14}** between measurements at **2.5 dB** below and **2.5 dB** above optimum power. This has led us to look for other biases that are sensitive to power. Since the upgraded servo computer code was installed, power shifts, if present, have been no larger than the Type A uncertainty. We attribute this change to an improvement in the way that the main frequency servo treats the optical transient following a cycle of the C-field servo. For biases linear in microwave power, such as Rabi pulling, Ramsey pulling, or rf sideband pulling, their effect on the frequency at **2.5 dB** below optimum power is just **0.39** times the observed frequency difference between measurements **2.5 dB** below optimum and **3.0 dB** above optimum. The ability to make this large change in microwave power gives us some leverage in looking for unknown power-dependent biases. However, microwave leakage has a different power dependence.

We have carried out many other parametric tests, some of which are described elsewhere in this paper. We have changed the modulation frequency, depth and

even the modulation sequence and duty cycles. We have changed the entire servo code. We have changed the laser power, polarization, beam size, pumping transition and zone in the beam tube. We have changed the number of lasers used (from a single laser for both pumping and detecting to separate lasers), the type of laser (from extended-cavity to DBR type), and the laser frequency-control system (from low-frequency modulation to an rf sideband technique). We have altered the atomic beam trajectory by the use of movable masks placed in the beam. We have altered the C-field strength and direction. We have changed the way the electronics are packaged and cabled. We have even changed the laboratory temperature and humidity. None of these parametric tests has yielded frequency changes significantly larger than their Type A uncertainty.

8. Summary and discussion

In the preceding sections we have described in detail the procedures we use to evaluate NIST-7. We have introduced the concept of an idealized frequency measurement (1) where all the biasing effects can be separated from pure, white noise effects and where the frequency of a reference oscillator can be expressed directly in terms of the SI second. We use this framework to guide our routine evaluations.

We evaluate all known biases by leveraged experiments that relate through a model some more easily measured quantity to the bias of interest. We have discussed the biases considered and the validity of the models used in their evaluation. This has allowed us to reduce the uncertainty on the individual biases below the Type A uncertainty of our frequency measurements.

Our approach has allowed us to reserve parametric measurements for model verification and the search for “hidden” biases. We vary all manner of conditions to include the type and set-up of the lasers, the cabling of the various components of the servo system, and the environment of the standard. If we have taken care of all significant biases, these variations should have no effect on the output frequency. While this philosophy for the operation of a primary standard places a great burden on the operator, we believe this is not a shortcoming of the technique. On the contrary, it represents the only technique we know to search for the unknown and often extremely subtle biases that exist in systems of this complexity.

In accordance with the requirements of (1) we monitor the noise residuals from all our frequency measurements for any indication of an unknown bias. We use the Allan variance as a routine check on individual frequency measurements. We have described other analytical techniques we have used to investigate the noise behaviour over different time intervals. During uncertainty evaluations, we never observe any indication of a noise type other than white.

Tables 3 and 4 summarize all the biases we have considered and their uncertainties. All uncertainties are quoted as 1σ .

Table 3. Relative frequency biases and their standard uncertainties.

Physical effect	$10^{15} \times$ Bias	$10^{15} \times$ Uncertainty
Second-order Doppler	≈ -300	1
Second-order Zeeman	$+10^5$	0.1
Cavity pulling	-6	0.6
Cavity phase (end-to-end)	± 615	0.7
Cavity phase (distributed)	-1.3	0.4
Black body	-20.4	0.3
Gravitation	179.9	0.1
Uncorrected biases	0	3.2
Combined Type B uncertainty		3.5

Table 4. Standard uncertainties for uncorrected biases.

Effect	$10^{15} \times$ Uncertainty
Magnetic field inhomogeneity	0.03
Rabi pulling	0.02
Ramsey pulling	0.002
Bloch-Siegert shift	0.3
Fluorescent light shift	0.5
Majorana transitions	< 1.3
Collisions	1.0
Beam flux variation	≤ 0.1
Microwave leakage	< 1.0
DC Stark shift	0.01
Spectral purity	0.1
Modulation synchronous effects	
Detector/demodulator	1.0
AM on laser	1.0
Switching transients	2.0
Combined Type B uncertainty	3.2

The major uncertainties in biases due to physical effects are in the second-order Doppler and end-to-end cavity phase biases. These uncertainties are limited by how well we can determine velocity distributions and compute the relevant velocity averages. They are a consequence of using a broad velocity distribution. However, our overall Type B uncertainty is dominated by uncertainties in electronic biases. This is a consequence of the unprecedented degree to which we split the Ramsey fringe: almost three parts in 10^7 .

These uncertainties have changed from our previous publications [5, 6]. The changes come from additional measurements made since 1994, upgrades for the lasers, electronics and software, and better understanding of some biases. The uncertainties for the end-to-end phase bias and for electronic biases have been greatly reduced. More biases have been considered.

Because of the indirect means we use to evaluate the biases, their associated uncertainties are expected

to be highly uncorrelated. For this reason, we use the root sum of squares of all the individual uncertainties as our total Type B uncertainty. The result is 3.5×10^{-15} . Combining this with our smallest Type A uncertainty of 2×10^{-15} for a routine evaluation, we arrive at a combined standard uncertainty (CSU) of 4×10^{-15} . This value represents our “best possible” performance, not our typical uncertainty. For many routine evaluations, especially those before MJD 51500, we did not achieve the uncertainties shown in Table 3 because some bias had not been adequately studied, or the standard was not operating under optimal conditions, e.g. only one active atomic beam was available.

Figure 12 shows the results of every reported evaluation of NIST-7. The data represent the departure d of the duration of the second measured by various primary standards from the realization of the SI second as calculated by the BIPM [63, 64]. These data include noise from the instability of TAI, of the primary standards, and from time transfer to the BIPM. From the NIST to the BIPM the transfer noise is about 1×10^{-15} over a 30-day interval [65]. The cavity-tuned hydrogen maser reference allows the few days encompassed by a NIST-7 evaluation to be extended to 30 days without significant additional uncertainty [65]. The change in level of CS2 in 1995 (MJD 49900) was due to inclusion of the black-body bias, a bias always included in NIST-7 evaluations. The change in level of all the standards between MJD 50400 and 50900 was due to the gradual steering of TAI to accommodate the black-body bias.

When we first started reporting to the BIPM in 1994, our frequency measurements rapidly changed by a total of 4×10^{-14} , larger than our then reported uncertainty of 1×10^{-14} . No single cause was ever pinpointed for this change. It occurred as our electronics were improved, external microwave leakage was eliminated, and our evaluation procedures were refined. Reduced atom-beam flux in late 1996 caused larger Type A uncertainties and reduced stability. The

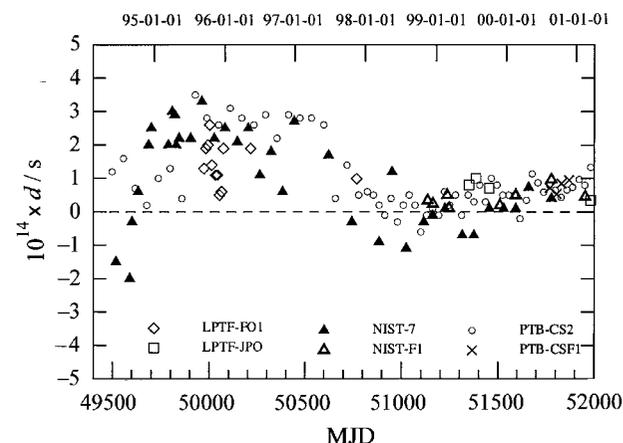


Figure 12. Comparison of NIST-7 and other primary standards with TAI. The data include transfer noise.

scatter in NIST-7 evaluations late in 1997 and early 1998 suggested that we then had an uncontrolled bias. Our evaluation techniques had also not yet matured to the level reported in this paper.

In the summer of 1998 a source of internal microwave leakage was closed. Improved servo-control was implemented in late 1999. Since then, the variance between TAI and NIST-7 has been smaller. The two low points in mid-1999 (MJD 51300-51400) and the high point in April 2000 (MJD 51661) were made with data in only one beam direction and have CSUs of 1.2×10^{-14} and 0.8×10^{-14} , respectively. But for two-beam evaluations our reported combined standard uncertainty has now been reduced to 5×10^{-15} , close to our estimate of the “best possible” uncertainty achievable with NIST-7. We undoubtedly still have unknown biases, perhaps related to the electronics or to the optical pumping. But the agreement of NIST-7 with recent evaluations of our in-house caesium fountain NIST-F1 [24] (also shown in Figure 12), suggests that any remaining unidentified biases are no larger than our reported combined standard uncertainty.

Acknowledgements. This has been a massive project extending over fifteen years. The input of many people has been vital to its success. It is impossible to mention them all. However, Len Cutler, Andrea De Marchi, Daniele Rovera and Fred Walls were exceptional for their technical input and insight. David Glaze, John Lowe and Lew Mullen were instrumental in the original design and assembly.

Note. Contribution of US government, not subject to copyright.

References

- Resolution 1, 13th Conférence Générale des Poids et Mesures, *Metrologia*, 1968, **4**, 41-45.
- Recommendation S1, BIPM Consultative Committee for Time and Frequency (CCTF), 14th Meeting, April 1999.
- Glaze D. J., Hellwig H., Allan D. W., Jarvis S. Jr, Wainwright A. E., *IEEE Trans. Instrum. Meas.*, 1974, **IM-23**, 489-501.
- Wineland D. J., Allan D. W., Glaze D. J., Hellwig H. W., Jarvis S. Jr, *IEEE Trans. Instrum. Meas.*, 1976, **IM-25**, 453-458.
- Lee W. D., Shirley J. H., Lowe J. P., Drullinger R. E., *IEEE Trans. Instrum. Meas.*, 1995, **IM-44**, 120-123.
- Drullinger R. E., Shirley J. H., Lee W. D., *Proc. 28th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, 1996, 255-263.
- Rovera G. D., de Clercq E., Clairon A., *IEEE Trans. Ultras. Ferroel. Freq. Contr.*, 1994, **UFFC-41**, 245-249.
- Hagimoto K., Ohshima S., Nakadan Y., Koga Y., *IEEE Trans. Instrum. Meas.*, 1999, **IM-48**, 496-499.
- Lee W. D., Drullinger R. E., Shirley J. H., Nelson C., Jennings D. A., Mullen L. O., Walls F. L., Parker T. E., Hasegawa A., Fukuda K., Kotake N., Kajita M.,

- Morikawa T., *Proc. IEEE International Frequency Control Symposium*, 1999, 62-65.
10. Bauch A., Fischer B., Heindorff T., Schröder R., *Metrologia*, 1998, **35**, 829-845.
 11. Bauch A., Heindorff T., Schröder R., Fischer B., *Metrologia*, 1996, **33**, 249-259.
 12. Recommendation S2, BIPM Consultative Committee for Time and Frequency (CCTF), 14th Meeting, April 1999.
 13. Drullinger R. E., Shirley J. H., Glaze D. J., Hollberg L. W., *Proc. 40th Symposium on Frequency Control*, 1986, 428-431.
 14. Drullinger R. E., Glaze D. J., Lowe J. P., Shirley J. H., *IEEE Trans. Instrum. Meas.*, 1991, **IM-40**, 162-164.
 15. Drullinger R. E., Shirley J. H., Lowe J. P., Glaze D. J., *IEEE Trans. Instrum. Meas.*, 1993, **IM-42**, 453-456.
 16. Ramsey N. F., *Molecular Beams*, Oxford, Clarendon Press, 1956, 466 p.
 17. De Marchi A., Shirley J., Glaze D. J., Drullinger R., *IEEE Trans. Instrum. Meas.*, 1988, **IM-37**, 185-190.
 18. Drullinger R. E., Glaze D. J., Sullivan D. B., *Proc. 39th Symposium on Frequency Control*, 1985, 13-17.
 19. Tanner C. E., Wieman C., *Phys. Rev. A*, 1988, **38**, 1616-1617.
 20. Theobald G., DeMarq N., Giordano V., Cerez P., *Opt. Comm.*, 1989, **71**, 256-262.
 21. Garcia Nava J. F., Walls F. L., Shirley J. H., Lee W. D., Delgado Aramburo M. C., *Proc. IEEE International Frequency Control Symposium*, 1996, 973-979.
 22. Lee W. D., Shirley J. H., Walls F. L., Drullinger R. E., *Proc. 49th Symposium on Frequency Control*, 1995, 113-117.
 23. Lee W. D., Shirley J. H., Drullinger R. E., *Proc. IEEE International Frequency Control Symposium*, 1997, 249-252.
 24. Jefferts S. R. et al., submitted to *Metrologia*.
 25. Giordano V., Pichon L., Boussert B., Cerez P., Blot P., *Proc. 8th European Frequency and Time Forum*, 1994, 359-367.
 26. Itano W. M., *Proc. 5th Symposium on Frequency Standards and Metrology* (Edited by J. C. Bergquist), Singapore, World Scientific, 1995, 503-504.
 27. Shirley J. H., *J. Appl. Phys.*, 1963, **34**, 783-788.
 28. Shirley J. H., *Phys. Rev.*, 1965, **138**, B979-B987.
 29. Vanier J., Audoin C., *The Quantum Physics of Atomic Frequency Standards*, Bristol/Philadelphia, Adam Hilger, 1989.
 30. Feynman R. P., Vernon F. L., Hellwarth R. W., *J. Appl. Phys.*, 1957, **28**, 49-52.
 31. Rabi I. I., *Phys. Rev.*, 1937, **51**, 652-654.
 32. De Marchi A., Rovera G. D., Premoli A., *Metrologia*, 1984, **20**, 37-47.
 33. Shirley J. H., *Proc. 43rd Symposium on Frequency Control*, 1989, 162-165.
 34. Shirley J. H., *IEEE Trans. Instrum. Meas.*, 1997, **IM-46**, 117-121.
 35. Mungall A. G., *Metrologia*, 1971, **7**, 49-56.
 36. Mungall A. G., *Metrologia*, 1972, **8**, 28-32.
 37. De Marchi A., Rovera G. D., Premoli A., *IEEE Trans. Ultras. Ferroel. Frequ. Contr.*, 1987, **UFFC-34**, 582-591.
 38. Barnes J. A. et al., *IEEE Trans. Instrum. Meas.*, 1971, **IM-20**, 105-120.
 39. Lee W. D., Shirley J. H., Drullinger R. E., *Proc. 48th Symposium on Frequency Control*, 1994, 658-661.
 40. Shirley J. H., Lee W. D., Rovera G. D., Drullinger R. E., *IEEE Trans. Instrum. Meas.*, 1995, **IM-44**, 136-139.
 41. De Marchi A., Drullinger R. E., *Metrologia*, 1987, **24**, 23-25.
 42. Weiss M. A., Ashby N., *Metrologia*, 2000, **37**, 715-717.
 43. Itano W. M., Lewis L. L., Wineland D. J., *Phys. Rev. A*, 1982, **25**, 1233-1235.
 44. Simon E., Laurent P., Clairon A., *Phys. Rev. A*, 1998, **57**, 436-439.
 45. Cutler L. S., Flory C. A., Giffard R. P., De Marchi A., *J. Appl. Phys.*, 1991, **69**, 2780-2792.
 46. De Marchi A., Rovera G. D., Drullinger R., Howe D. A., *Proc. 39th Symposium on Frequency Control*, 1985, 3-7.
 47. Bloch F., Siegert A., *Phys. Rev.*, 1940, **57**, 522-527.
 48. Stevenson A. F., *Phys. Rev.*, 1940, **58**, 1061-1067.
 49. Shirley J. H., *Proc. 39th Symposium on Frequency Control*, 1985, 22-23.
 50. Makdissi A., Berthet J., de Clercq E., *IEEE Trans. Ultras. Ferroel. Frequ. Contr.*, 2000, **UFFC-47**, 461-465.
 51. Ohshima S., Nakadan Y., Ikegami T., Koga Y., *IEEE Trans. Instrum. Meas.*, 1991, **IM-40**, 1003-1007.
 52. Bauch A., Schröder R., *Annalen der Physik*, 1993, **2**, 421-449.
 53. De Marchi A., *Metrologia*, 1982, **18**, 103-116.
 54. Kajita M., *Z. Phys. D*, 1997, **41**, 3-8.
 55. Lee W. D., Lowe J. P., Shirley J. H., Drullinger R. E., *Proc. 8th European Frequency and Time Forum*, 1994, 513-516.
 56. Rossi F., Opat G. I., *J. Phys. D*, 1992, **25**, 1349-1353.
 57. Audoin C., Jardino M., Cutler L. S., Lacey R. F., *IEEE Trans. Instrum. Meas.*, 1978, **IM-27**, 325-329.
 58. Nelson C. W., Walls F. L., Sicardi M., *Proc. IEEE International Frequency Control Symposium*, 1994, 567-571.
 59. Parker T. E., *Proc. 13th EFTF and 1999 IEEE International Freq. Control Symp.*, 1999, 173-176.
 60. Greenhall C. A., Dick G. J., *IEEE Trans. Ultras. Ferroel. Frequ. Contr.*, 2000, **UFFC-47**, 1593-1599 and references therein.
 61. Lesage P., *IEEE Trans. Instrum. Meas.*, 1983, **IM-32**, 204-207.
 62. Available on the Internet at http://www.bipm.org/enus/5_Scientific/c_time/time ftp.html
 63. Thomas C., *Proc. 11th European Frequency and Time Forum*, 1997, 283-289.
 64. Thomas C., Azoubib J., *Metrologia*, 1996, **33**, 227-240.
 65. Parker T. E., Howe D. A., Weiss M., *Proc. IEEE International Frequency Control Symposium*, 1998, 265-272; Parker T. E., *ibid.*, 2001, 57-62.

An index of publications by the NIST Time and Frequency Division is available on the Internet at <http://www.bldrdoc.gov/timefreq/index.html>

Received on 12 February 2001 and in revised form on 9 May 2001.