Chapter 8
The Rubidium Clock

8.1 The Reference Hyperfine Transition

Of the atomic clocks, or more appropriately, frequency/time standards, since their accuracy and sophistication, not to mention their cost, places them far above any ordinary keepers of time, the rubidium clock has the distinction of being the most compact, and therefore the most portable. Rugged versions of the rubidium standard have long been developed for shipboard use as well as for tactical military and missile-borne applications.

The rubidium standard is based on the resonance at microwave frequency of the free rubidium atom between a pair of its quantum states whose separation in energy is due to the electron–nuclear hyperfine interaction. Its compactness is a result of confining the rubidium vapor in a small absorption cell filled with a noble gas to act as a buffer, as mentioned in the last chapter. While this method of confining the atoms of rubidium in order to lengthen their free interaction time with the applied resonant field has been very successful, there are unfortunately residual effects on the frequency of resonance due to the collisions with the buffer gas and to the pumping light itself that disqualify it as an absolute standard. Nevertheless, its general adoption for a variety of applications attests to its usefulness as a secondary standard.

In order to appreciate how the reference transition was chosen from among the many possible resonances observable in the ground state of Rb, we must examine the way in which the energies of the various magnetic substates may depend on the environment, since the resonance frequencies are directly determined by the difference in energy between quantum states. The most important environmental factor is the magnetic field; as earthbound beings immersed in the earth's magnetic field and surrounded by man-made magnetic fields from machinery, etc., it would require elaborate special shielding or active field-cancellation to reduce all static and time-varying magnetic fields. Fortunately, a particular choice of a resonant transition exists, whose frequency is very much less sensitive to the

F.G. Major, The Quantum Beat
© Springer Science+Business Media, LLC 2007

magnetic field than all the others. To see this we must consider how the energy of the atom in the various hyperfine states varies with magnetic field intensity. We recall that Rb^{87} has a nuclear spin $I = 3/2$ and electron angular momentum in the ground state $J = 1/2$, leading to total angular momentum (hyperfine) states with $F = 2$ and $F = 1$. Each of these states comprises a set of substates labeled by their magnetic quantum number $m_F = 2, 1, 0, -1, -2$, and $m_F = 1, 0, -1$, which give the projections (or components) of the angular momentum, and hence magnetic moment, along an assumed magnetic field direction. In the presence of a magnetic field, this can be only an approximate way to describe the states, since the angular momentum F obtained by adding (vectorially) the nuclear and electron spins will no longer be strictly constant in magnitude and direction, due to the torques exerted on the particles by the magnetic field. In trying to compare the relative strength of the "coupling" between the spins with their tendency to precess independently around the magnetic field direction, the appropriate measure is the amount of energy that would be required on the one hand to turn one spin relative to the other as compared with turning it with respect to the field.

The coupling energy of the spins is the difference in energy between the $F = 2$ and $F = 1$ hyperfine states in zero magnetic field, and it is precisely the transition between these states that gives rise to the sharp microwave resonance used as the frequency reference. Now, the energy of coupling of the spins to the magnetic field is simply the Zeeman energy, which we have already encountered; it is given by $E_m = -\mu_{||}B$, where $\mu_{||}$ is the component of the magnetic moment of the electron along the magnetic field direction. We are now ready to express the condition on the strength of the magnetic field for the representation in terms of F and m_F to be a good approximation; we must have $E_{F=2} - E_{F=1} \gg \mu_{||}B$.

For magnetic fields weak enough to satisfy this condition, the combination of two spins acts as one, with a single angular momentum F precessing around the magnetic field; however, the magnetic moment associated with this differs from the electron moment by only the small contribution from the nucleus. Thus it is approximately as if we had a single particle with the ratio of magnetic to angular momentum smaller than a free electron in the ratio (1/2):F, and a correspondingly slower precession around the magnetic field. Thus in the limit of a vanishingly small magnetic field, the energy of the magnetic substates is simply $E_m = (m/F)\mu B$; that is, the plots of E_m versus B start from $B = 0$ with a different slope for each m, as shown in Figure 8.1.

8.2 The Breit–Rabi Formula

The behavior of E_m as the magnetic field intensity is increased from zero requires an exact quantum treatment; the result is referred to as the *Breit–Rabi formula*, which can be written as follows:

$$E(m_F) = -\frac{E_{hfs}}{2(2I+1)} - \frac{\mu_I B_0}{I}m_F \pm \frac{E_{hfs}}{2}\sqrt{1 + \frac{4m_F}{2I+1}x + x^2}, \qquad 8.1$$

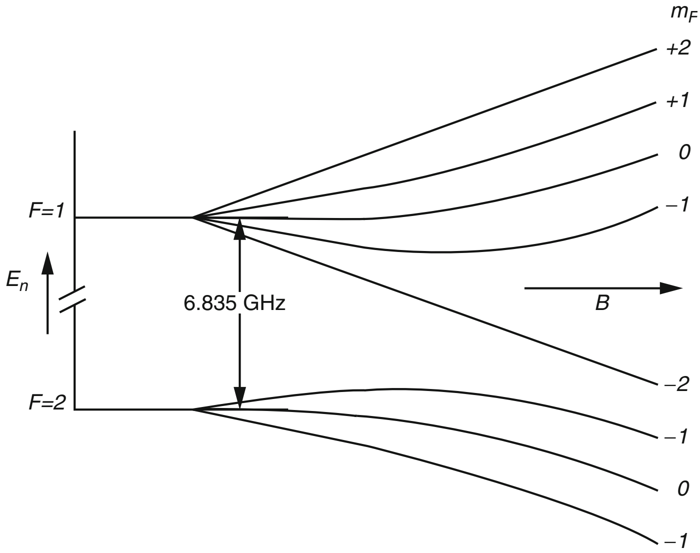

Figure 8.1 The energies of the magnetic hyperfine substates of the ground state in Rb^{87} as a function of applied magnetic field intensity

where the plus and minus signs refer to the upper and lower hyperfine state respectively, and $x \approx g_J \mu_B B_0 / E_{hfs}$ is the ratio of the Zeeman to the hyperfine energy splitting, and B_0 is the strength of the magnetic field.

There are two features of this solution of great importance to the operation of the rubidium standard: First, the plots of the energies of atoms in levels ($F = 2$, $m_F = 0$) and ($F = 1$, $m_F = 0$) versus the magnetic field start in a horizontal direction before they start curving gently, which means that to a first approximation in x the energy does not change if the magnetic field departs slightly from zero; second, the difference in energy between consecutive Zeeman sublevels near $B = 0$ is proportional to the magnetic field.

Since, as we have already stated, it is difficult in practice to totally shield out the perturbations of a magnetic field or even variations in its intensity over all points in the rubidium absorption cell, the field-insensitive transition between the two $m_F = 0$ levels belonging to $F = 1$ and $F = 2$ is used as the standard. Even the presence of a small inhomogeneity in the magnetic field, which would result in a given rubidium atom in motion "seeing" a variable magnetic field, would cause little spectral broadening of the resonance between these states. The same cannot be said obviously of the "field-dependent" transitions between other substates. Thus the resonance between the $m_F = 0$ states is much sharper than the other Zeeman transitions, a fact of obvious importance for a frequency standard.

The field independence of the energy of the $m_F = 0$ states is true only to a first-order approximation in the immediate vicinity of zero field; beyond that we must use the exact Breit–Rabi formula to calculate the frequency of the transition between these states. A second order approximation in x for that frequency derived from that formula yields

$$\nu = \nu_0 + \frac{(g_J\mu_B - g_I\mu_n)^2}{2h^2\nu_0}B^2, \qquad 8.2$$

where g_J and g_I are respectively the g-factors of the atomic electron and nucleus, numbers that are measures of the strengths of their magnetic moments in the given states. They specify the moments in terms of the fundamental units, the Bohr magneton μ_B and the analogous nuclear magneton μ_n, defined as the classical magnetic moment of a particle having the charge and mass of a proton. This provides a convenient way to make a fine adjustment to the frequency of the standard: simply vary the current in a magnetic field-producing coil provided for the purpose. As a secondary standard, it is necessary to set the field at such a value that the time scale generated agrees with the atomic time scale, defined in terms of the primary standard. Even after the initial calibration against the primary standard, a readjustment may be necessary after some length of time because of possible long-term drift in the resonance frequency. Furthermore, to estimate the size of the field correction, which is proportional to B^2, we have ready at hand the "field-dependent" Zeeman transitions $m_2 - m_1 = \pm 1$, whose frequency gives directly the magnitude of the magnetic field.

8.3 Optical Pumping of Hyperfine Populations

Recall that in our discussion of magnetic resonance in Chapter 6 we argued that in order to be able to observe a magnetic transition between two states, there must be a difference in the populations in those states. This is ultimately because the inherent probabilities (per atom) per unit time for absorption and stimulated emission of a quantum of radiation are identical, and unlike transitions in the optical region of the spectrum, the probability for spontaneous transitions is extremely small; the result is that no net global exchange of energy is observed unless the number of atoms in the lower state differs from the number in the upper state. Here we wish to observe the transition between the $(F = 1, m_F = 0)$ and $(F = 2, m_F = 0)$ states near zero magnetic field intensity. There are a number of possible ways of achieving this using optical resonance; the choice generally adopted reflects the inevitable concern for commercial viability, which is ultimately a question of performance versus cost. Before describing in detail the method that has been widely exploited commercially, we can get a broader perspective by first considering alternative, but more complex, approaches.

The first among them is simply to carry out the usual Kastler optical pumping with circularly polarized light directed along the magnetic field axis, thereby ideally putting all the atoms in one of the extreme m_F-states, that is, $m_F = +2$ or -2, depending on the sense of the circular polarization. We are assuming here, as we did in the last chapter, that the spectrum of the pumping light is limited to resonance with the transition to the upper $P_{1/2}$ electronic state but otherwise is broad enough in frequency to satisfy energy conservation for all hyperfine transitions that satisfy the angular momentum selection rules. Once we have a large proportion of the atoms in, say, the $m_F = +2$ state, we can apply a high-frequency magnetic field resonant with transitions among the magnetic substates in the $F = 2$ hyperfine state, in the manner described in the last chapter. The desired population difference between the $m_F = 0$ substates can be achieved under suitable conditions, the most important of which are first, that the atoms be sufficiently free of perturbations; second, that the high-frequency field be strong and uniform; and last, that the static magnetic field be sufficiently uniform. The global effect on the atoms can be pictured classically as a magnetized gyroscope whose axis precesses around the static magnetic field, sweeping out ever wider cones. If the high-frequency magnetic field is on for exactly the short interval it takes the angle of the cone to reach 90°, a so-called 90° pulse, then the axis of the gyroscope precesses in the plane perpendicular to the static field, and the projection of its angular momentum along the field is zero. In quantum terms this is described as having put the atoms in a (linear) superposition of substates with different m_F in which the desired $m_F = 0$ substate has the largest amplitude. Having a significant fraction of the atoms in the $(F = 2, m_F = 0)$ substate, a fraction far greater than would ideally be present in the $(F = 1, m_F = 0)$ substate, meets the first requirement for observing transitions between them.

A second and equally critical requirement is the ability to detect the occurrence of transitions; this can be met in principle by the inverse process of applying a 90° pulse of the opposite phase, or a 270° pulse of the same phase, to complete a full circle, bringing the global moment back into alignment with the static magnetic field. If nothing perturbs the atoms in the interval between the two pulses, the atoms would ideally return to their original state, namely, the nonabsorbing $m_F = 2$ into which they were pumped, and the amount of pumping light scattered or absorbed would be the same as it was. However, if in the interval between pulses a resonant microwave field causes transitions to the $(F = 1, m_F = 0)$ hyperfine state, then they do not all return to their nonabsorbing state, and the amount of pumping light absorbed/scattered will be increased. This change in the interaction between the atoms and the pumping light can be used to monitor the transitions and their resonant dependence on the frequency of the microwave field.

Another technique, which avoids the practical complexity of pulsed operation, is a variant of Kastler optical pumping, in which the circularly polarized beam is directed not along the magnetic field direction, but perpendicular to it, and is therefore called *transverse pumping*. In order to simplify the explanation of this technique, let us assume that the transverse light beam consists of a regular succession of powerful flashes and that the magnetic field is extremely weak. Under these conditions each flash produces a global magnetic polarization in the direction of the beam, that is, perpendicular to the field, which then causes it to precess continuously around the field axis passing periodically through its original direction. Now, if the interval between flashes is adjusted so that each flash coincides in time with the passing of the polarization vector through its original direction, then the polarization is reinforced and will build up to a significant degree. Again, having the polarized atoms precessing predominantly perpendicular to the field implies a preponderance of population in the $m_F = 0$ state, essentially the same state as was produced by a 90° pulse. In the actual implementation of the transverse pumping technique it is not necessary to pulse the light source. Instead, a high-speed electro-optic modulator can be used to impose a harmonic oscillation in the transmitted intensity at the frequency of precession of the atoms in the given magnetic field. If during the pumping process the desired microwave transition is resonantly induced between the hyperfine states having $m_F = 0$, then the distribution of populations of atoms in the different m_F states changes in the direction of increasing those in the absorbing $F = 1$ substates. It will be recalled that the optical pumping process leads to a preponderance of atoms in substates that by reason of the selection rules are unable to absorb light from the pumping beam. Hence by imposing a different distribution with the resonant microwave transitions, the amount of pumping light scattered by the atoms will increase, providing a way of monitoring the resonance. As with the previous technique then, the desired resonance is observed by monitoring the transmitted light intensity as the microwave frequency is swept through resonance; a dip in the intensity of the pumping light transmitted through the absorption cell signals a resonance. It is interesting to note in passing that the original use of a circularly polarized light beam perpendicular to the magnetic field was first introduced by Hans Dehmelt (Dehmelt,1957) as the inverse process to the foregoing: It was to modulate the intensity of the beam by interaction with a precessing global polarization induced by a resonant high-frequency magnetic field, such as would be used in the 90° pulse, acting on polarization produced by an axial beam. This modulation occurs at the precession frequency and is a direct measure of the (static) magnetic field. Since frequency is measurable with high precision, this has been exploited commercially as an atomic magnetometer of great sensitivity and precision.

8.4 Optical Hyperfine Pumping: Use of an Isotopic Filter

We will now direct our attention to the principle of operation actually implemented in commercial Rb standards; it is called *hyperfine pumping*. Instead of relying on quantum selection rules governing transitions between states of different angular momentum, it is really based on selection of transitions according to the conservation of energy. We focus on the spectrum of the pumping light rather than its polarization; transitions will occur only if the energy and therefore wavelength of the photons in the beam equals the energy difference between the initial and final states. The method therefore relies on having a pumping light source whose spectrum overlaps only one of the two hyperfine components in the resonance optical spectrum of Rb, components that arise from transitions whose initial states are either the $F = 1$ or $F = 2$ hyperfine state. Having a spectrum overlapping only one component, the light from the source can be absorbed by atoms in only one of these states. However, once in the optically excited state, the atoms will spontaneously re-emit photons to both hyperfine states of the electronic ground state, independently of how they came to be in the excited state, and therefore ideally they would all be pumped into the nonabsorbing hyperfine state.

As mentioned in connection with light sources for Kastler pumping, laser sources properly stabilized would be ideal were it not for the added complexity that that would entail. It might be thought that we should be able simply to filter out one of the hyperfine components of the optical resonance line in the spectrum of a Rb vapor lamp; unfortunately, the difference in wavelength between the two components is so small that it would be difficult, if not impossible, using the sharpest type of optical filter available, the interference filter, to separate them without a great loss of intensity.

The original experiments on hyperfine pumping predate lasers, and a suitable light source was achieved through a fortuitous coincidence in the hyperfine structure of the optical resonance spectra of the two isotopes of rubidium, Rb^{85} and Rb^{87}. The difference in nuclear structure and mass of the two isotopes leads to a slight relative shift in their spectra, called, not surprisingly, the *isotope shift*. It happens that one of the two hyperfine components in the Rb^{85} spectrum nearly coincides with the corresponding component in the Rb^{87} spectrum, while the others are well separated, as shown in Figure 8.2. Thus starting with a rubidium vapor lamp filled with enriched Rb^{87} isotope, whose output contains both hyperfine lines, we can partially remove one of them by passing the light through a cell containing enriched Rb^{85} vapor, which will absorb out of the beam (and re-emit in all directions) the line coincident with the Rb^{85} hyperfine component just mentioned. The match in wavelengths can be improved by a process that, however, we shall see is detrimental to the long-term stability of the standard: It is the so-called pressure

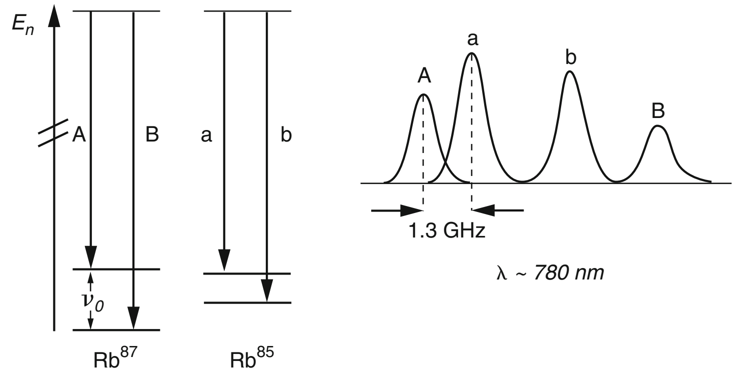

Figure 8.2 The hyperfine structure of the dominant emission lines in Rb^{85} and Rb^{87}

broadening and shift in spectral lines caused by collisions between the Rb atoms themselves and with others, principally the atoms of the noble gas introduced as a buffer. The direction of the shift, whether to higher or lower wavelengths, depends on which noble gas is used, as does its sensitivity to temperature fluctuations. By using a mixture of two noble gases, which alone would produce opposing shifts, it is possible to choose their proportion in the Rb^{85} vapor absorption cell so as to reduce the temperature dependence as well as improve the wavelength match.

The earliest successful observations of resonance between hyperfine states in the alkali atoms by optical means date from 1958. Of these, the one of particular interest, because of its adoption for further commercial development, was published by Bender, Beaty, and Chi (Bender et al., 1958), in which the use of the isotope filter in Rb was introduced. Their experimental arrangement is shown in Figure 8.3. The optical hyperfine pumping source was a Rb spectral lamp whose strongest emission occurs at the two "resonance" lines in the red part of the spectrum at wavelengths $\lambda = 780\,\mathrm{nm}$ and $795\,\mathrm{nm}$ ($1\,\mathrm{nm} = 1$ nanometer $= 10^{-9}$ meter). These correspond, in terms of the states between which the transitions occur, to the strong emission lines in a sodium vapor lamp, giving it its familiar yellow color. They arise from radiative transitions between the first two excited electronic states and the ground state, forming a fine structure "doublet." The two hyperfine states with $F = 1$ and $F = 2$, into which the first excited electronic state is split, are much closer in energy than the two corresponding hyperfine states in the ground state, so that under the usual degree of resolution the spectrum appears to have each member of the doublet split into two hyperfine components rather than four. An ordinary Rb

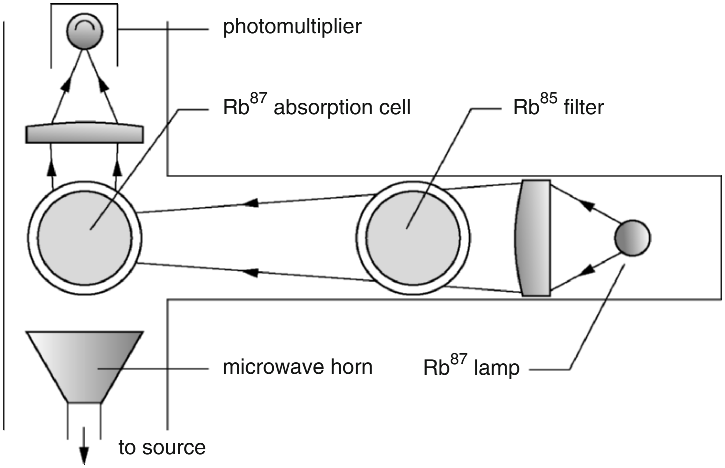

Figure 8.3 The experiment of Bender, Beaty, and Chi on optical hyperfine pumping of Rb^{87}

vapor lamp, containing a natural mixture of the two isotopes, will therefore emit a spectrum in which each line of the doublet consists of four hyperfine components, two from each isotope. The hyperfine separation in Rb^{85} is about half that in Rb^{87}, with the lower components of the two isotopes much closer in frequency than the upper, as Figure 8.2 attempts to make clear. In the actual experiment, the light from an enriched Rb^{87} spectral lamp is passed through a filter cell containing enriched Rb^{85} vapor and typically 10^4 Pa of argon buffer gas, whose presence broadens the spectral lines and somewhat shifts their centers in a direction to enhance the differential filtering of the two hyperfine lines of Rb^{87}. Ideally, if the filter cell contained only Rb^{85}, it would be almost transparent to the upper component in the output spectrum of the Rb^{87} lamp, while strongly scattering the other component, so that the transmitted light satisfies the basic inequality of intensities required for hyperfine pumping. The repeated cycle of absorption and re-emission of this light by the Rb^{87} in the resonance absorption cell will pump the atoms into the upper hyperfine state of the electronic ground state, thereby reducing the number left in the lower absorbing state. This has the effect of decreasing the amount of pumping light scattered. If an applied microwave magnetic field causes transitions between the hyperfine states, tending to equalize their populations and thus increasing the number in the lower absorbing state, then more of the pumping light is scattered out of the beam, and the transmitted intensity drops, signaling the occurrence of resonance.

8.5 The Use of Buffer Gases

In those early experiments, the absorption and filter cells were made of relatively large ($\approx$500 ml) Pyrex bulbs, which were cleaned and baked according to standard vacuum practice, and a small quantity of pure metallic Rb distilled into each bulb. The element Rb, like the other alkalis, is very chemically reactive with air and water; it must be handled either in an inert atmosphere or under vacuum. The cells are back-filled with pure gas before being sealed from the vacuum system.

Experiments carried out with different species of gases at different pressures showed that the hyperfine frequency is shifted in proportion to the pressure, being raised by the light gases—hydrogen, helium, neon, and nitrogen—while for the larger atoms—argon, krypton, xenon, and methane—the frequency is lowered. The linear pressure dependence is to be expected, at a fixed temperature, if the shifts are due to pairs of atoms colliding. There is also an important temperature dependence, however, which is more complicated to predict. An important practical application of these findings is to the reduction of the sensitivity of the frequency to temperature, since some gases cause an increase in frequency with temperature, while others have the opposite effect. It has been found that a mixture of about 12% Ne and 88% Ar gives a temperature dependence of about -10 hertz per degree C at a pressure of about 5.320×10^3 Pa (1 mm Hg $\approx$ 133 Pa). The sharpest resonance seen for the hyperfine transition at 6.8347... GHz (1 GHz $= 10^9$ Hz) was only about 20 Hz! This is a Q-value of about 300 million; compare this to the best quartz crystal currently available, which may reach a Q-value of 1 million.

The effects of a buffer gas on spectral line shapes and center positions are the large-scale average manifestation of interaction between the Rb atoms and the noble gas atoms; during the collision we have in effect a transient Rb–noble gas "molecule." Under the rarefied conditions obtaining here there would be little error in assuming what may be called the *binary collision approximation*. In this, it is assumed that all the particles have negligible interaction except very briefly during relatively infrequent encounters when the particles come within typical molecular dimensions of each other. This radically simplifies the problem of predicting the effect of the presence of the buffer gas by permitting the problem to be separated into two tractable parts: first, the collision of just two particles under general conditions, and second, the statistical problem of finding the observable averages over a large number of such collisions. Another circumstance that allows an important simplification in the analysis of the collision process is that the relative velocity of the Rb and noble gas atoms at the temperatures under consideration is very much smaller than the speed with which the electrons in the outer structure of the atoms can adjust to the changing internuclear distance. This means that the colliding pair can be thought of as quasi-static at different distances apart, and the energy of the quantum electron state for the two-atom system can be regarded as potential energy in computing the change in kinetic energy of the two particles. The potential energy for the Rb–noble gas collisions, in common with most binary atomic collisions and in a broader sense all matter, corresponds to an attractive force as the particles first approach each other; but as their electronic structures start to interpenetrate, the force turns repulsive and they fly apart. It turns out that during the initial attractive part of their trajectory, the distortion of the electronic state of the Rb atom is accompanied by a reduction in the electron–nuclear interaction, that is, a "red" shift (to lower frequency) in the hyperfine frequency separation, while the repulsive part has the opposite effect. In the heavier noble gases the attractive force is of longer range and leads to a net red shift, while for the lighter gases, He and Ne, it is a blue shift (to higher frequency) because the short-range repulsive force dominates. The length of time a typical collision lasts is extremely small, as can easily be verified: The average relative velocity of the particles due to thermal agitation is on the order of 10^4 meters per second, while the range of interatomic force is typically 10^{-8} meter, so that the time is $t = d/V = 10^{-12}$ second. On the other hand, the average time between collisions at the typical gas pressure of 1000 Pa is about 10^{-7} second, or about 100,000 times the duration of a collision. This confirms that the impact approximation is indeed valid for the assumed conditions. Of course, if very much higher buffer gas pressures are used, then the approximation would become invalid and the prediction of the pressure shifts would be very much more difficult. The temperature dependence of the frequency shifts has to do with degree of mutual penetration of the colliding particles; the higher their thermal kinetic energy, the more violent the collisions.

The susceptibility of the resonance frequency, which is to be our standard, to temperature, pressure, and nature of background gases detracts from its accuracy and reproducibility. However, the presence of the buffer gas not only performs the essential function of lengthening the interaction time between the atoms and the resonant field, but also limits broadening of the resonant response by another phenomenon, namely the Doppler effect.

The frequency width at half maximum of the optical resonance line in rubidium is about 700 MHz, which is about 1/10 the hyperfine splitting of the ground state. On the other hand, we can easily verify that for the microwave transition the conditions for the Dicke effect are well satisfied; thus the wavelength of microwave resonance at around 6.8 GHz is $\lambda = c/\nu$, or $\lambda \approx 4.3$ cm, and the average distance between collisions at pressures on the order of 10^4 Pa is no more than 0.01 cm, or about 1/400 of the wavelength. Finally, on the subject of the Doppler effect we should note that if the resonant microwave field is applied in the form of an advancing wave, as was done in the experiments of Bender, Beaty, and Chi, the observed resonance will exhibit a small Doppler shift due to a general drift of Rb atoms across the absorption cell. This arises from the fact that Rb reacts chemically with the glass surfaces of the cell, and atoms continually diffuse from their source, a droplet of liquid Rb, towards the walls of the cell. A good deal of effort has been devoted in the past to finding a suitably inert coating for the inner surfaces of the cell, not only to reduce this chemical reaction, but indeed to dispose of the need for a buffer gas altogether. Special aluminosilicate glazes have been developed by the lamp industry to coat the inner surfaces of sodium lamps used for street lighting, lamps that operate at much higher temperatures than the Rb cell, in order to prevent the sodium vapor from chemically attacking the glass and eventually turning it black. The ideal of a cell with surfaces totally unreactive with Rb, which may therefore be "dry filled," that is, not requiring a liquid droplet to maintain the vapor density, has never been achieved. Nevertheless, wax coatings made of high molecular weight paraffins were shown by H. Robinson et al. around 1957 to be highly successful in preventing a randomization of Rb spin direction when the atom collides with the coated surface. Their use instead of a buffer gas to increase the free interaction time between the Rb atoms and the resonant field offers, however, no particular advantage, since like the buffer gases, these coatings also cause frequency shifts. In any event, whether the glass surfaces are coated or not, there will be long-term chemical reaction with the Rb, resulting in the slow evolution of gases, which will cause the resonance frequency to drift. The use of sapphire and other exotic materials to solve this problem continues to be investigated.

8.6 Light Shifts in the Reference Frequency

An equally serious but far more subtle phenomenon that affects the resonant microwave frequency is associated with the pumping light itself; this complex effect is labeled simply the *light shift*. It was anticipated theoretically by

Cohen-Tannoudji and Barrat in 1961 and was soon observed in Kastler's laboratory in the radio-frequency spectrum of Hg^{199}. Its detection in studies of microwave resonance in Rb and Cs was first published in the same year by Arditi and Carver. Although the shift is exceedingly small—its discovery in itself representing no mean accomplishment—yet in the context of a frequency standard it is significant. In addition to the obvious dependence of the shift on the intensity of the light, it also depends on the detailed spectrum of the light, particularly its position in relation to the absorption spectrum of the Rb atoms. In trying to understand the physical origin of light shifts we must distinguish between two types of quantum transitions involved in the optical pumping cycle: real transitions and the so-called *virtual* transitions. A transition is virtual if the probability of the atom being in the final state returns to zero when the external perturbation is removed rather than stay there, as in a real transition. In the case of Rb subjected to the optical pumping cycle, real transitions would take the atoms up to the first excited electronic state, where they would stay were it not for another mechanism, spontaneous emission, by which they re-emit photons and return to the ground state. At the same time there are virtual transitions in which the electric field of the light wave distorts the electron cloud, and as long as that field is there the atoms are in a quantum state that in terms of the undisturbed Rb stationary states can only be described as a *linear superposition* of them. Once the perturbing electric field is removed, the atom returns to its initial state. This distortion in the electron distribution manifests itself in a shift in the position of lines in the optical spectrum called the (quadratic) *AC Stark effect*, and it depends in general on E^2, where E is the amplitude of the electric field in the light wave. Because of the quadratic dependence on E, the effect does not average to zero for an oscillating optical field that swings symmetrically through positive and negative values about zero. Unfortunately, the energy shift of the quantum states is not the same for the initial and final hyperfine states of the microwave transition we are interested in, resulting in a light shift in the frequency of that transition. Since this source of change in our standard frequency is affected by the many complex factors that determine the detailed spectrum of the light source and absorbing Rb atoms, it seriously detracts from the quality of the Rb standard.

8.7 Rubidium Frequency Control of Quartz Oscillator

We will now take up the subject of the electronic configuration of clocks controlled by the Rb resonance. There are two ways in which the microwave resonance can be used: first, in a passive mode as a resonator or frequency discriminator, and second, in an active mode as an oscillator (maser) generating a signal at the standard frequency. The passive Rb standard is the one that has been commercially developed and is in general use; we will therefore treat it in this chapter, leaving the Rb maser for a later chapter.

As with other standards using an atomic resonator in a passive mode, notably the cesium standard, which we shall study in the next chapter, the resonant response of the atoms to an external microwave field must be monitored. Specifically, it must be possible to discern whether the frequency of the field is below, above, or precisely at the center of the resonance curve. This may be accomplished by a slow periodic modulation of the frequency (or phase) of the microwave field over a small portion of the resonance line profile, as shown in Figure 8.4. We recall that if the modulation is slow enough, it is legitimate to think of the frequency as assuming continuously all frequencies between the limits of the modulation, so that the optical signal will vary in step with the modulation on the side of the resonance curve with a positive slope, and will vary in the opposite direction on the side with the negative slope. If the modulation occurs symmetrically about the center of the resonance, then since the optical signal falls whether the field frequency swings in the positive direction or the negative, the optical signal will oscillate at twice the frequency of modulation. If the modulation is exactly centered on the peak of the resonance, then the optical signal will have no Fourier component at the modulation frequency, but only one at double that frequency.

For the ultimate purpose of controlling the frequency of the microwave field so that it remains locked to the peak of the resonance, we need to derive from the optical signals described above a voltage that can serve as an error signal indicating whether the frequency of the field is too high or too low. This requires a circuit that can selectively amplify signals at the modulation frequency and be sensitive to the relative phase of these signals with respect to the modulating signal. Such a phase-sensitive amplifier is called a *lock-in amplifier*, which in a sense correlates an incoming signal with a reference frequency signal; if the two are in phase, a positive output voltage is given: On the other hand, if they are of opposite phase,

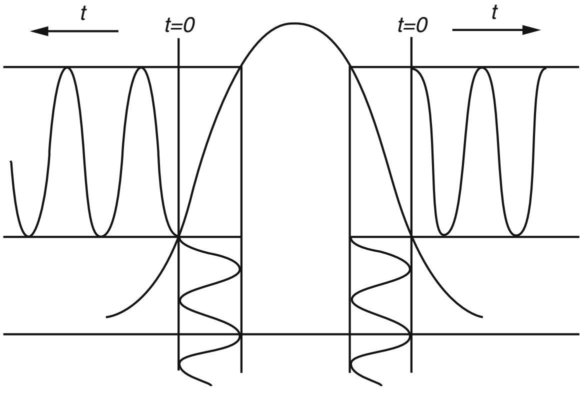

Figure 8.4 The reversal of detector output phase for microwave frequency above and below resonance

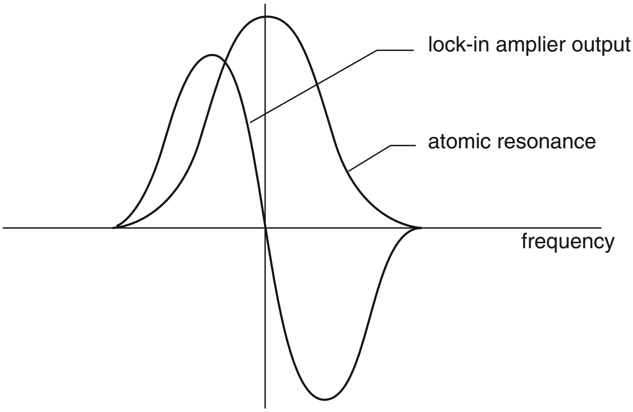

Figure 8.5 The output of the lock-in amplifier as the frequency is swept through resonance

we get a negative output voltage. Moreover, if they are not exactly the same frequency, the output will oscillate as the two signals go in and out of phase; in this case the average over a sufficiently long time will be zero. With such a lock-in amplifier we can obtain just the desired error signal for our control circuit; we simply use the signal producing the frequency modulation of the microwave field as the reference, and the optical signal as our input to the lock-in amplifier. Its output as the center frequency of the field is very slowly scanned across the resonance frequency will resemble the plot shown in Figure 8.5 A graph of that form is usually called a dispersion curve, and in the sense used in calculus, it is the derivative function of the bell-shaped absorption curve. We see that the output is indeed negative when the frequency is too high, positive when too low, and zero when at the peak. This is precisely what we need as an error signal in a feedback control loop that seeks to make the error zero by controlling the center frequency of the field.

The field is derived ultimately from a high-quality quartz-controlled oscillator operating typically at 5 MHz (see Figure 8.6). The desired microwave frequency is produced by a frequency synthesizer, which, starting with the 5 MHz oscillation as reference, generates signals at multiples and submultiples of that frequency and then by deriving other signals at the sum or difference of various harmonics, ultimately yields an output signal whose frequency can be preset in fine increments on a front panel keyboard.

These arithmetical operations on the frequencies of signals are realized through the use of nonlinear solid-state devices, which can act as harmonic generators and frequency mixers. Throughout these operations phase relationships are preserved so that the output is coherent with the stable 5 MHz reference signal. Coherence here simply means that the outputs of two identical synthesizers, sharing the same reference signal but set at different frequencies, can produce a stable "beat," that is, a pure signal whose frequency is the difference between those of the two outputs. Clearly, if the frequency of the reference is changed, the frequency of the output of

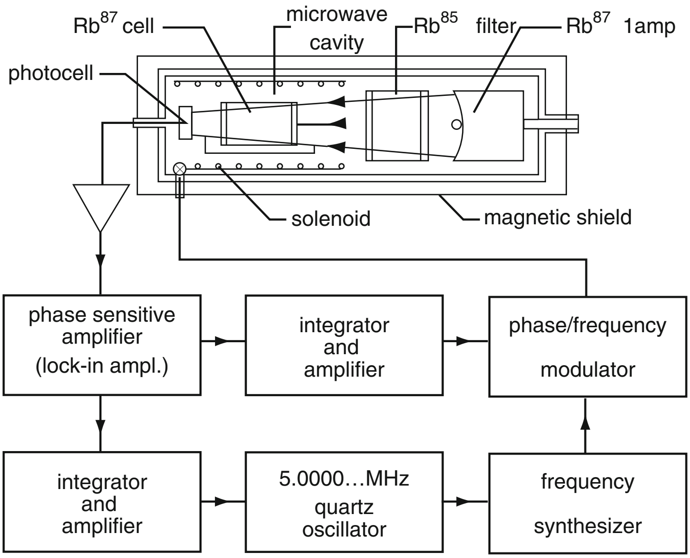

Figure 8.6 Block diagram of basic rubidium frequency standard design

the synthesizer will change accordingly. Therefore, a voltage-sensitive element in the quartz crystal oscillator circuit is provided to control its frequency and hence that of the microwave field applied to the atoms. If that frequency is too low, the resulting positive error signal must cause the oscillator frequency to rise steadily in order to reduce the error. This requires the voltage appearing on the frequency-controlling element, which is derived from the output of the lock-in amplifier, to increase steadily in magnitude in the proper direction. The circuit that converts a steady (DC) voltage into a linearly increasing one is an *integrator*, which must be included in the feedback loop for stable operation. As the error signal approaches zero at the peak of the resonance, the output of the integrator tends to become constant just at the value to keep the error at zero. Needless to say, a high degree of stability in the operating voltages of the integrator is critical; any drift in DC levels would cause frequency offsets from the atomic resonance, degrading the performance of the standard. Just how closely the frequency is held to the true center of the atomic resonance depends on a large number of factors, some of which are of a fundamental nature, while others are a matter of the performance characteristics of particular devices and the circuits around them.

8.8 Frequency Stability of the Rubidium Standard

We have already mentioned a number of physical phenomena that affect the frequency of the atomic resonance; these are known systematic sources of error, as distinguished from uncontrollable random fluctuations. Among the latter are the various types of electrical noise discussed in a previous chapter. Because of the high order of multiplication of frequency in going from that of the quartz oscillator at 5 MHz to the microwave frequency at 6,800 MHz, any residual fluctuations in the phase of the quartz oscillator are greatly magnified. Therefore, attaining the highest phase stability in the microwave field and hence the sharpest spectrum puts a great burden on the spectral purity and low noise of the quartz oscillator. It is the availability of high-performance quartz oscillators, with extremely high Q-values and low output noise that has contributed immensely to the success of these atomic standards.

In discussing the accuracy and stability of any type of standard, questions must be addressed that would not arise for ordinary instruments. When we use an ordinary voltmeter, for example, we assume that its calibration and accuracy are traceable ultimately to some acceptable standard. But if the standard be in doubt what then? This question is really relevant only to the cesium standard, which has been elevated to the status of primary time standard. However, as an atomic resonance-based system, the rubidium clock qualifies, for many applications, as a secondary standard and as such, absolute accuracy is not expected of it; its frequency must be set by reference to a primary standard. But how is one to know whether the primary standard is drifting? This question lies at the heart of what is expected of a standard: Standards are not supposed to drift! The pragmatic answer is to have a large collection of embodiments of the standard all purporting to display a unit of time in accordance with its atomic definition. To the extent that there is agreement among the members of this collection, we can have confidence in their accuracy and stability.

The accepted method of specifying the stability of frequency standards, useful particularly for relatively long-term performance, is, as we saw in a previous chapter, in terms of the *Allan variance* of phase or frequency plotted as a function of the sampling time over which that quantity is measured. We recall that this analysis presumes that the condition of stationarity is satisfied, and therefore any long-term drift in the data must first be separated out. We also noted previously that some of the fundamental types of noise can be accurately modeled as having a Fourier spectrum that has a simple power-law dependence on frequency. These power laws translate into equally simple dependence of the Allan variance $\sigma(\tau)$ on the time interval τ used in its measurement. Thus for the important flicker noise it can be shown that σ is independent of the length of the interval, whereas for white frequency noise σ falls as $1/\tau^{1/2}$. Now, for circuits at ordinary temperatures operating in the radio-frequency range, thermal (Johnson) noise is very nearly "white" (the same power density at all frequencies), so that since this is a universal source of noise, we frequently see a plot of σ versus τ exhibit the $1/\tau^{1/2}$ characteristic of this type of noise, at least up to a certain point, after which flicker noise becomes dominant and the graph flattens out.

Figure 8.7 shows plots of the Allan variance for a typical Rb standard together with, for the sake of comparison, several other types of atomic standards we shall be dealing with in succeeding chapters. It is not unusual for a Rb standard to

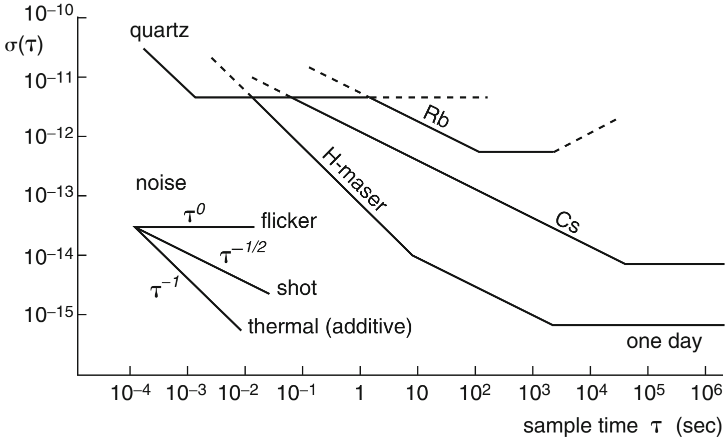

Figure 8.7 Allan variance plotted versus sampling time for different types of frequency standards

have σ bottom out at around 10^{-12} for time intervals longer than 100 seconds, after systematic drifts have been separated out. This represents an error of about 30 millionths of a second in a year!

8.9 The Miniaturization of Atomic Clocks

The optically pumped alkali vapor gas-diffusion cell resonator lends itself admirably to miniaturization, particularly since the development of semiconductor lasers that are tunable and can be stabilized on the rubidium or cesium optical resonance wavelengths while operating at room temperature. This development obviated the need for the UHF driven lamp and isotope filter of the conventional rubidium standard. Coupled with microelectronic integrated circuitry this has enabled atomic clocks to be built no larger than a walnut! One such clock is shown in Figure 8.8.

The fundamental consequences of size reduction are an increased frequency of wall collisions or the need for a higher collision frequency with the molecules of the buffer gas. In either case the undesirable shifts in the reference transition frequency are aggravated, and long-term stability is expected to suffer. The other important consideration is the size of a resonant microwave cavity. For Rb^{87} and Cs^{133} the reference microwave wavelengths are about 4.4 cm and 3.3 cm respectively. Since Rb has fewer magnetic substates than Cs, a larger fraction of Rb atoms can contribute to the signal arising from transitions between one particular pair. In either case the cavity must be “loaded” with a low-loss dielectric material to lower its resonant frequency as its size is made smaller. It has been shown (I. Liberman, 1992), for

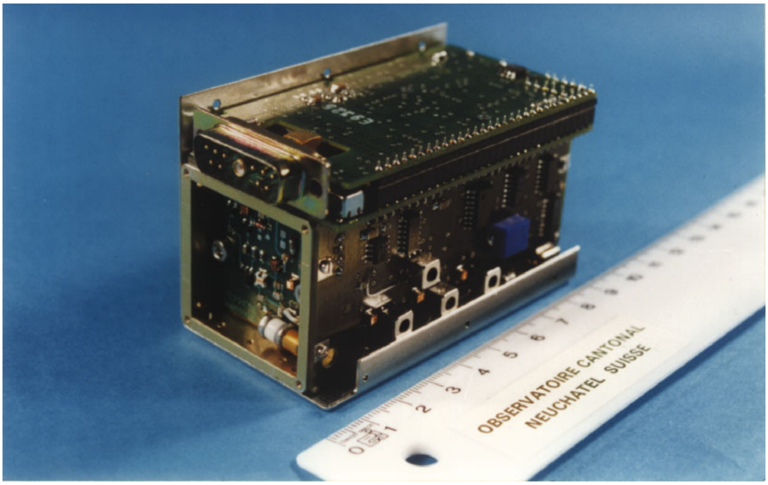

Figure 8.8 Small "industrial" Rb clock developed by the Neuchâtel observatory (Rochat, 1994)

example, that a miniature standard based on a cesium gas cell not exceeding 4 mm in diameter and 18 mm long, operating under relatively high vapor pressure, so that Cs–Cs collisions are dominant in limiting the free lifetime of the Cs states, would have theoretically a short-term stability $\sigma(\tau)$ on the order of $5 \times 10^{-12}\tau^{-1/2}$ where τ is in seconds.