Eur. Phys. J. D **7**, 449–460 (1999)

THE EUROPEAN PHYSICAL JOURNAL D

EDP Sciences

A diode-laser optical frequency standard based on laser-cooled Ca atoms: Sub-kilohertz spectroscopy by optical shelving detection

C.W. Oates[a], F. Bondu, R.W. Fox, and L. Hollberg

National Institute of Standards and Technology, Mail Stop 847.10, 325 Broadway Boulder, CO 80303, USA

Received 2 November 1998

Abstract. We report an optical frequency standard at 657 nm based on laser-cooled/trapped Ca atoms. The system consists of a novel, compact magneto-optic trap which uses 50 mW of frequency-doubled diode laser light at 423 nm and can trap $> 10^7$ Ca atoms in 20 ms. High resolution spectroscopy on this atomic sample using the narrow 657 nm intercombination line resolves linewidths (FWHM) as narrow as 400 Hz, the natural linewidth of the transition. The spectroscopic signal-to-noise ratio is enhanced by an order of magnitude with the implementation of a "shelving" detection scheme on the 423 nm transition. Our present apparatus achieves a fractional frequency instability of 5×10^{-14} in 1 s with a potential atom shot-noise-limited performance of $10^{-16}\,\tau^{-1/2}$ and excellent prospects for high accuracy.

PACS. 32.80.Pj Optical cooling of atoms; trapping – 42.55.Px Semiconductor lasers; laser diodes – 42.62.Fi Laser spectroscopy

1 Introduction

The work described here focuses on the unique features of laser cooling of the alkaline earth atom Ca, as we prepare this system for use as an optical frequency/wavelength standard. The ultimate goal of our work is the development of a high performance and convenient optical frequency reference. One of the selling points of the Ca-based system is that the cooling and clock wavelengths are accessible with diode laser systems, leading to the possibility of a simple, compact, and ultimately transportable apparatus. However, most of the seminal work thus far has used dye lasers for both trapping and spectroscopy. In this paper we present details of a diode-laser-based Ca magneto-optic trap (MOT) and sub-kilohertz diode-laser spectroscopy on the 657 nm transition. In the course of this work we have developed new trapping and spectroscopic strategies, which are tailored to the idiosyncrasies of the alkaline earth atom. Those features of laser cooling peculiar to these atoms have profound implications for both our trapping and spectroscopic approaches. Furthermore, we have exploited the level structure of our effectively two-electron atom by implementing a "shelving" detection technique [1,2]. This method, akin to methods used in trapped-ion studies, substantially improves our signal-to-noise ratio.

It is now evident that the use of laser-cooled atoms and ions will lead to orders of magnitude improvement in stability and accuracy for the next generation of frequency standards [3,4]. Enhanced stability results primarily from the increased atom-light interaction time provided by slow atoms, which leads to narrower transition linewidths. Increased accuracy results from narrower lines and the reduction of critical velocity-dependent systematic frequency offsets, such as the second-order Doppler shift. In the microwave region, cold atoms are a great success in the Cs atomic fountain [4]. In the optical domain, improvements due to laser cooling may prove even more significant. Indeed, over the past few years there has been excellent progress in the development of the next generation of optical frequency/wavelength standards, due in large part to the use of laser-trapped atoms and ions. There are now several systems poised to move us from the several kilohertz uncertainty of the familiar I_2-stabilized HeNe system toward the hertz level and even beyond [1,5–9].

The trapped-ion candidates can achieve extremely narrow linewidths with high stability, and they have excellent prospects for extremely high accuracy [8,9]. Neutral atom/molecule candidates have the potential for extremely high stability due to the large number of participating atoms/molecules, and they should achieve good accuracy as well. Due to their narrow linewidths and insensitivity to external perturbations, the intercombination lines of the alkaline earth atoms are among the most promising and practical neutral candidates. In fact, high resolution spectroscopy has already been demonstrated on the lowest-lying intercombination lines of Mg [1], Ca [5],

[a] e-mail: oates@boulder.nist.gov

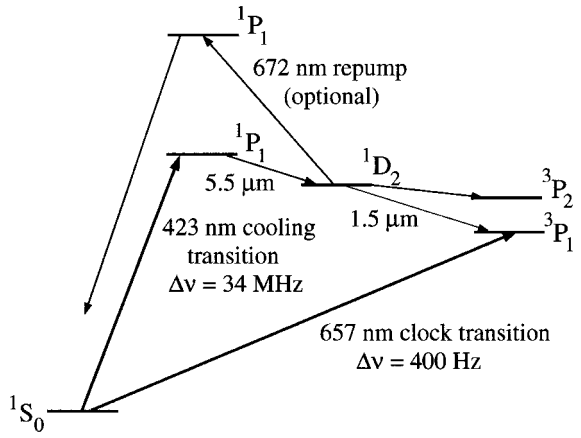

Fig. 1. Relevant energy levels for neutral ^{40}Ca.

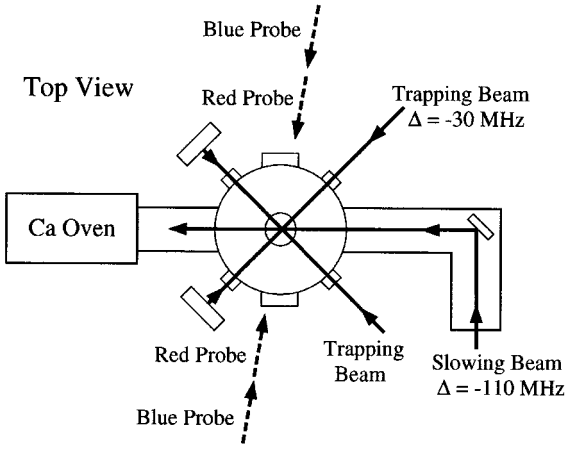

Fig. 2. Apparatus used for trapping Ca. Omitted for clarity are the vertical trapping beams and magnetic field gradient coils. See text for details.

Ba [10], and Sr [11]. For a frequency reference based on one of these transitions, the uncertainty for a thermal atom beam system can be as small as 500 Hz (dominated by the uncertainty in the correction for the second order Doppler shift). With laser-cooled atoms, however, preliminary investigations of the Mg and Ca systems have yielded estimates of a potential uncertainty of < 1 Hz, though much work still remains to reach this level [1,5]. Progress with cold atoms has been steady and recently several groups, including our own, have resolved sub-kilohertz linewidths (line $Q \sim 10^{12}$) on alkaline earth atom transitions [12–14]. Such well-defined frequency markers are useful not only as standards, but also for atomic interferometry and tests of fundamental physics.

The 657 nm $^1S_0\,(m = 0) \rightarrow {}^3P_1(m = 0)$ transition in Ca is particularly attractive due to its narrow natural linewidth (400 Hz) and convenient laser wavelengths for spectroscopy and cooling (see the level diagram in Fig. 1). For the development of our optical frequency reference the use of laser cooled/trapped samples is essential to provide the extended interaction time needed for sub-kilohertz spectroscopy and to reduce systematic effects. The pioneering work on high resolution spectroscopy of this transition was conducted in 1979 by Barger and collaborators, who were able to resolve linewidths as narrow as 6 kHz (FWHM) [15]. They used a thermal Ca beam with what was then a new spectroscopic technique, saturated absorption with optical Ramsey fringes [16–18]. Over the past 15 years workers at Physikalisch Technische Bundesanstalt (PTB) have advanced the technology and performed a variety of experiments on this transition as part of its development for use as an optical frequency standard [5]. Their work has included high resolution spectroscopy on laser-trapped Ca and a measurement of the absolute frequency of the 657 nm transition [19]. Their current uncertainty for this measurement is only 120 Hz, making this the most accurately known optical atomic transition frequency. Their work has led to the addition of this transition by Comité Consultatif pour la Définition du Mètre (CCDM) to the list of those recommended for the realization of the meter [20]. There are now several groups worldwide developing optical frequency/length references based on this transition.

2 Apparatus

2.1 A compact magneto-optic trap for neutral Ca

As we set out to construct a relatively compact MOT for Ca, we were faced with two main challenges. One was to construct a diode-laser-based system to generate sufficient light power at 423 nm for excitation of the $^1S_0 \rightarrow {}^1P_1$ cooling transition (see the following section). A second was to find a simple MOT design, which would work efficiently for Ca. Unfortunately, the simplest form of atom trap, which catches slow atoms from the background in a vapor cell [21], is not easily applied to alkaline earth atoms such as Ca. For adequate vapor pressure, we would need to heat the whole trapping apparatus to $> 400\,^\circ$C, which is possible but presents numerous experimental difficulties [22]. For this reason, the first Ca MOTs were loaded with atoms from Zeeman-slowed atomic beams [5,23]. While these traps performed well, they were large and required strong magnetic fields. We decided to try instead a modified version of a trap design which was first demonstrated for Li at Stanford [24] (and subsequently for Ca at PTB [12] and NIST [25]). This design foregoes the use of the Zeeman slower and relies instead on catching only the slow atoms from the atomic beam. To compensate the reduced velocity capture range, we place the trap as close to the atomic source as possible, so as to have a large atom flux through the trapping region.

In our present version (see Fig. 2) the trap center is located 13 cm from the oven. The oven, which we operate at $\sim 600\,^\circ$C, has a nozzle (1 mm diameter, 1 cm long) which serves as the only beam-collimating aperture in the system. The cylindrical vacuum chamber of the trapping region (10 cm in diameter) has eight windows in addition to entry and exit ports for the atomic beam.

Four of these windows are set in the horizontal plane at 45° relative to the atomic beam. These 1.5 cm diameter windows are anti-reflection-coated at 423 nm for the horizontal trapping beams. There is a pair of larger (4 cm diameter) horizontally oriented windows (perpendicular to the atomic beam) which are anti-reflection-coated for 657 nm and 423 nm to transmit our blue and red probe beams. A second pair of large, dual-wavelength-coated windows are oriented vertically (normal to the atomic beam). These transmit the vertical trapping beams and fluorescence from the trapped sample, which we collect and send to a photomultiplier tube (PMT). This arrangement limits the fluorescence detection coverage to 1% of 4π sr, but some imminent modifications will increase this by nearly an order of magnitude. A single ion pump (20 l/s) keeps the background pressure below 1.3×10^{-5} Pa (10^{-7} torr) during oven operation. The four large optical quality windows were mounted on the vacuum system using the Kasevich modified-copper-gasket method [26].

Another feature which differentiates the alkaline-earth MOT from the more widely used alkali MOT is the need for much larger magnetic field gradients. This results from the broader cooling transition linewidth, which is typically five times that of an alkali atom. Since the required magnetic field gradient approximately scales with the cooling transition linewidth [27], larger gradients are required for efficient trapping. For Ca we typically work with gradients ranging from 6 to 10 mT/cm. To generate these large gradients we send ~ 120 A through a pair of 10 cm diameter, water-cooled, anti-Helmholtz coils, which each consist of 16 turns of 0.6 cm diameter copper tubing. During the spectroscopic measurement cycle it is desirable to be able to chop the current in $< 100\,\mu$s, so we have constructed a transistor chopper for this purpose. In practice we rarely use this chopper (in part due to poor reliability, mechanical vibrations, and eddy currents), because we have learned that we can do the majority of our spectroscopy of the $m = 0 \rightarrow m = 0$ transition in the presence of the gradient field. Our trapped atom sample is horizontally displaced from zero field as a result of radiation pressure from the slowing beam and sees a ~ 0.5 mT field in a horizontal direction parallel to the atomic beam propagation. We thus orient the polarization of our red and blue probe beams horizontally to excite the magnetically insensitive transitions. We can measure the magnitude of the magnetic field (to $< 10\,\mu$T) as well as the gradient by rotating the red probe polarization to 45° and viewing all three of the $^1S_0 \rightarrow {}^3P_1$ components. Under present conditions, the gradient leads to less than 20 Hz of broadening on the $\Delta m = 0$ transition through the second-order Zeeman shift (assuming our measured value of 61(4) Hz/(mT^2) for the second-order Zeeman shift coefficient).

Before dividing our 423 nm light into the various trapping beams, we send it through two acousto-optic modulators (AOMs) to facilitate fast and complete turn-off. After division each beam passes through a quarter-wave plate in order to produce the circular polarization required by the trap. With about 10 mW per retroreflected horizontal beam and 2 mW for the retroreflected vertical beam, we can trap 10^6 atoms in 20 ms for a trap detuning of -30 MHz. The trap lifetime is limited to ~ 20 ms by optical pumping to the $^1D_2$ state (see Fig. 1), a feature shared by all alkaline earth atoms except for Mg. From this state an atom decays 25% of the time to the extremely long-lived $^3P_2$ level, in which an atom simply falls out of the trap. Otherwise an atom will decay to the $^3P_1$ level and back to the ground state on a time scale (3.5 ms) which gives a good probability for recapture. Since the 20 ms trap lifetime is much shorter than the several hundred milliseconds time we measure for collisions with thermal atoms to eject trapped atoms, it is not necessary to shield the trapped atomic sample from the atomic beam [28]. As we will describe in the Section 2.3, we have also been able to increase the lifetime through optical pumping out of the D state, thus increasing the number of atoms in our trap.

Another approach for increasing the equilibrium number of atoms is to increase the loading rate by accessing a larger portion of the atomic beam's velocity distribution. One way to increase the velocity capture range is to add a frequency-shifted sideband (or sidebands) to the trapping light [12,24]. We increase the loading rate by adding a slowing beam tuned 110 MHz below resonance, which is aligned counter-propagating to the atomic beam. Interestingly, we find the greatest increase in trapped atom number when we transfer most of the power from the trapping beams (still detuned 30 MHz to red of the transition) to the slowing beam. Our present version uses ~ 12 mW in the 0.8 cm diameter slowing beam and ~ 1 mW in each 1 cm diameter trapping beam. Even without a special Zeeman slower, this beam still slows enough atoms in 13 cm that our loading rate increases by an order of magnitude. In fact, though they are far from optimal for slowing the atomic beam, our large diameter trapping coils act as a crude Zeeman slower. The evidence for this is that we see an improvement ($\times 2$) in the loading rate when we switch the polarization of our slowing beam from linear to the correct circular. Conversely, we see no additional cooling for the opposite circular polarization.

With the small size of our trap, we cannot simply send the slowing beam directly through a window facing the oven, as the window becomes coated with Ca within a week and no longer transmits light. Instead we send the slowing beam through a window perpendicular to the atomic beam and reflect it off a mirror set at 45° inside the vacuum system (see Fig. 2). While this mirror is also susceptible to coating with Ca from the atomic beam, the metallic Ca-coated version of this mirror still gives adequate reflection ($\sim 40\%$ after three months of fairly regular use) at 423 nm for good performance of the slowing beam.

Perhaps the most critical difference between the laser cooling of alkalis and alkaline earths is the resultant temperature. Since the most abundant alkaline earth isotopes have no nuclear spin, their singlet ground states have no Zeeman structure, thus precluding the usual sub-Doppler-cooling mechanisms. Furthermore, the large linewidth of the cooling transition leads to a large Doppler-cooling

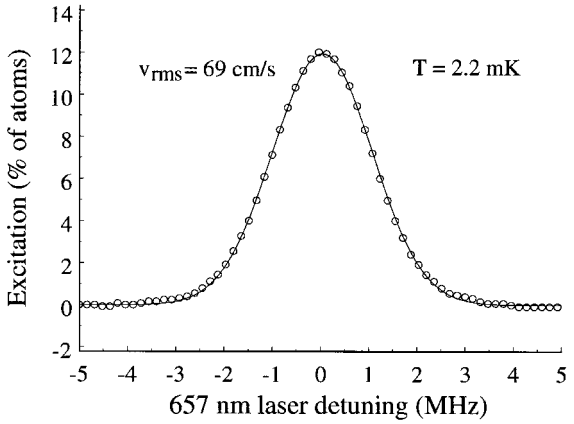

Fig. 3. Cold atom velocity distribution determined by the excitation induced by a single 2 μs red probe pulse. Note that 1 MHz = 0.6 m/s at the probe wavelength of 657 nm.

limit, which for Ca is 0.8 mK, corresponding to a root-mean-squared velocity ($v_{\rm rms}$) of 42 cm/s. We can readily measure the velocity distributions of our cold atoms using the first-order Doppler shift on the narrow red transition. In Figure 3 we show a typical velocity distribution, measured for a trapping laser detuning of -28 MHz. This lineshape closely resembles a thermal distribution with a $v_{\rm rms}$ of 69 cm/s, corresponding to a temperature of 2.2 mK. The resultant atomic temperature is actually a function of the intensity and detuning of the trapping laser light [5], but manipulation of these parameters has not yet enabled us to reach the Doppler limit. These comparatively warm temperatures for alkaline atom cooling have important implications for our spectroscopic investigations.

2.2 423 nm diode laser-based system

One of the biggest technical challenges in building a diode-based MOT for Ca is the generation of single-frequency, tunable radiation at 423 nm with sufficient power ($>$ 20 mW). Since there are no suitable 423 nm semiconductor sources yet available, we double the frequency of high-power light at 846 nm using a nonlinear $KNbO_3$ crystal (see Fig. 4).

The high-power, frequency-controlled light at 846 nm comes from a semiconductor master oscillator-power amplifier (MOPA) that consists of an extended-cavity diode laser (ECDL) and tapered amplifier. The ECDL uses a diffraction grating in the grazing-incidence configuration to provide frequency-selective optical feedback to the laser (Hitachi HL8318E [29]). Tunability of our extended-cavity diode laser is achieved through a PZT-controlled feedback mirror and to a lesser degree, the laser current. We stabilize the frequency of the master laser by locking it to an environmentally isolated Fabry-Perot cavity, whose spacer is fabricated from ultra-low expansion glass. This cavity rests on a v-block inside an evacuated chamber, which is passively isolated both thermally and acoustically. The long-term and thermally driven drifts of our cavity are then limited to a few megahertz per day, which has minimal effect on our trapping performance. Furthermore, we compensate these small drifts manually with a double-passed AOM at 80 MHz, which allows tuning of the trapping laser frequency relative to a cavity fringe which lies near the Ca trapping frequency. Finally, to increase our 846 nm output power, we send the 20 mW output from our ECDL to a tapered amplifier. This amplifier can produce 500 mW of power, but we operate it at 250 mW (with 1350 mA injection current) to avoid thermal effects in the doubling crystal.

This high-power light is then sent to our nonlinear frequency-doubling crystal, a 12 mm piece of $KNbO_3$, which is used in a resonant optical build-up cavity. The crystal facets are anti-reflection coated ($R < 0.25\%$). We achieve a light build-up factor of about 30 with a 4.5% (transmission) input coupler, which was chosen for this input power and conversion efficiency. With the aid of a pair of beam-shaping cylindrical lenses after the tapered amplifier, we are able to couple the light to the doubling cavity with 85% efficiency. The cavity is kept in resonance with the input light by dithering its length and feeding back to a PZT-controlled cavity mirror. The sub-freezing phase-matching temperature ($-10\,^\circ$C) requires a flow of dry N_2 or O_2 over the crystal to prevent condensation and icing on the crystal faces. Under optimal conditions, the system has produced 70 mW of blue light, although we usually work closer to 50 mW, due to better thermal stability of the crystal at this power. Furthermore, the intensity of this second-harmonic light is stabilized to 0.1% with a servo system which controls the tapered amplifier current. This doubling system has run reliably for six months now, while a previous system (which used an injection-locked 150 mW diode laser instead of an amplifier) ran for more than two years with only minor maintenance.

2.3 Increasing the trap lifetime through optical pumping from the $^1D_2$ state

In Section 2.1 we described how optical pumping to the $^3P_2$ through the $^1D_2$ state leads to a trap loss mechanism which limits the trap lifetime to $\sim$ 20 ms (depending on trap parameters). If we want longer trap lifetimes, we can add an ECDL at 672 nm to pump atoms from the $^1D_2$ state to a higher $^1P_1$ state (see level diagram in Fig. 1), intercepting the atoms before they have a chance to decay to the $^3P_{1,2}$ states. From this higher $^1P_1$ state, the atoms have several available decay channels, with the most probable decay simply back to the $^1D_2$ state. After several cycles, however, they preferentially decay back to the atom ground state. This pumping time is on the order of microseconds, so the atoms are then readily recaptured by the trapping light. Ultimately this process is limited by (i) decay from the higher $^1P_1$ state to longer-lived triplet levels and (ii) generation of dark states in the singlet D level on this $J = 2 \rightarrow J = 1$ transition. The generation of dark states, however, might be less of a problem in the presence of the strong magnetic field gradient.

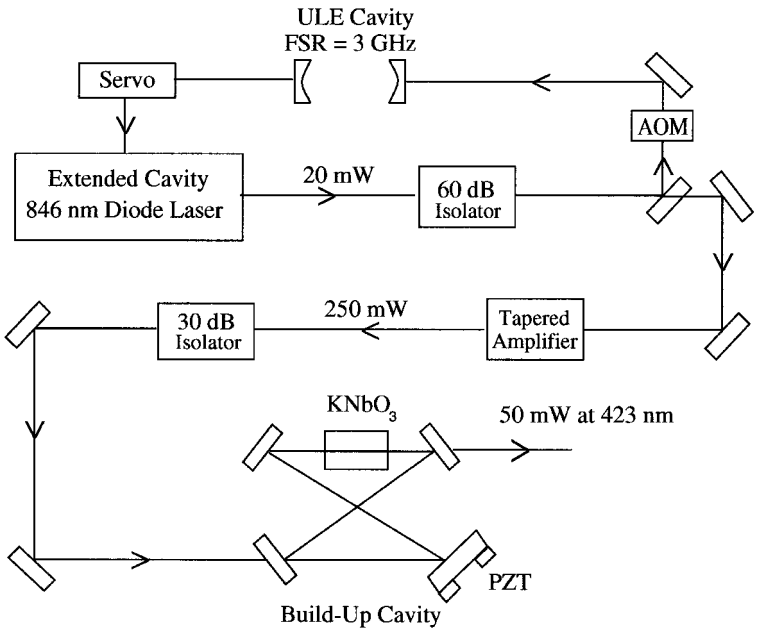

Fig. 4. Apparatus used for generating the 50 mW of tunable, frequency stabilized 423 nm light for trapping and probing.

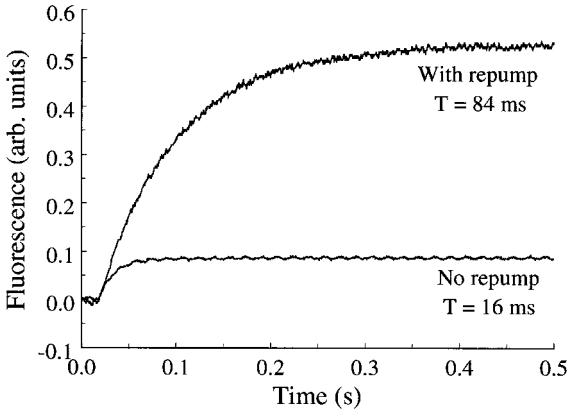

Fig. 5. Increase in trap lifetime and trap fluorescence due to the presence of 672 nm repumping light. We see here an increase in trap lifetime of more than a factor of 5, while a measurement taken at lower oven temperature showed a ninefold lifetime enhancement.

In a preliminary measurement of the repumping efficiency, we used 1.5 mW of 672 nm light (retroreflected) and saw the trap lifetime extended by as much as a factor of 9. This leads to a similar increase in the steady state number in the trap. This result was insensitive to polarization, perhaps due to the presence of the magnetic field gradient. With the optical pumping lifetime greatly extended, the trap lifetime was then limited by collisions with fast Ca atoms. We have yet to repeat such measurements with an atomic beam block in place. Figure 5 shows the increase in trap lifetime when we add the 672 nm repumping light. This demonstrates that with the addition of a single diode laser we can significantly increase the trap lifetime, which could improve the signal-to-noise ratio of the frequency reference. This increase in trap lifetime could also permit studies of collisions between cold alkaline earth atoms, a topic of interest to theorists due to simplicity of the atomic system [22]. These collisions are of interest to frequency metrologists as well, since for a cold atomic sample of sufficiently high density, collisions can cause significant systematic shifts [30].

2.4 657 nm laser system for high resolution spectroscopy

At the heart of our spectroscopic investigation lies our 657 nm diode laser system (see Fig. 6). The requirements for this system are extreme in that we need frequency stability suitable for sub-kilohertz spectroscopy and high laser power (> 5 mW, ideally 100 mW or more) with a good spatial mode. For diode lasers, these requirements push the limits of available technology. The fast laser linewidth (on a timescale < 1 ms) needs to be well below our chosen resolution (for example, 400 Hz), because the effective laser linewidth leads to decay in signal contrast by a factor of $\exp(-2\pi\Delta\nu T)$ [1], where $\Delta\nu$ is the laser linewidth and T is the time between spectroscopic pulses in the Ramsey technique (which will be described in more detail later). For a 400 Hz resolution ($T = 628\,\mu$s), we would then need a fast laser linewidth of < 50 Hz in order to have a $< 18\%$ decay of our signal contrast.

To achieve such a narrow linewidth we use two stages of stabilization for our anti-reflection coated, 657 nm Toshiba 9321 diode laser [29]. The first stage consists of a 657 nm ECLD in the Littman configuration, which reduces the fast laser linewidth to several hundred kilohertz, still far, however, from our requirements. For the second stage

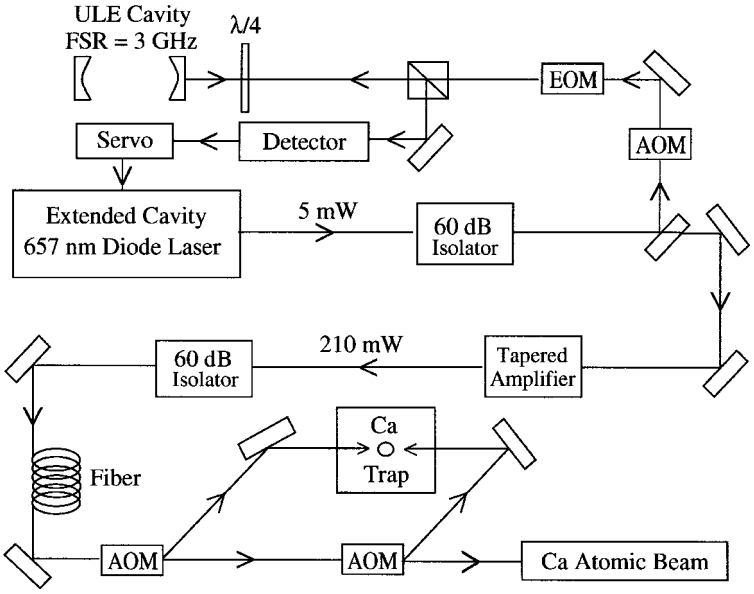

Fig. 6. Apparatus used for generating tunable, frequency-stabilized 657 nm probe light.

of stabilization, we split off 100 μW of the 4 mW laser output and couple it to the TEM_{00} mode of a high-finesse Fabry-Perot cavity (linewidth = 50 kHz). We then lock the laser frequency tightly to a suitable cavity resonance with a servo system which controls the laser current and PZT-controlled mirror. An error signal with a high signal-to-noise ratio is produced with the Pound-Drever-Hall technique [31], for which we generate FM sidebands with an electro-optic modulator driven at 13.3 MHz. Fast frequency fluctuations are corrected with the laser current, while slow corrections are sent to the PZT, so that the DC value of the laser current remains constant over long times (the laser will stay locked for a day or more). This servo combination yields a unity gain bandwidth of over 3 MHz, which enables us to reduce the linewidth to less than 50 Hz, determined from the frequency fluctuations measured relative to the cavity. While we think there is still room for improvement in the servo design, this stabilization has proven adequate for the work described here.

With our laser locked tightly to the cavity, our long term ($>$ 1 ms) laser frequency fluctuations are a result of changes in the cavity length due primarily to thermal and acoustic fluctuations. To reduce the sensitivity to environmental fluctuations, we have constructed our cavity mirror spacer out of ultra-low expansion glass and placed it in a chamber held under vacuum by a 2 l/s ion pump. For stable, kinematic mounting and small thermal conductivity, the cavity sits on sapphire balls resting in grooves cut into a v-block made of Invar. For the first layer of acoustic and thermal isolation, we put the vacuum chamber inside a temperature-stabilized metal box lined with lead foam. For further isolation we put this metal box inside of a second box with thermal insulation in between. Floor vibrations are also a serious issue since they can couple to the cavity and write frequency noise onto the laser. We use active feedback on the airlegs of the optical table to isolate the table from vibrations at frequencies above the vertical and horizontal table resonances of $\sim$ 0.7 Hz. The active control regulates the air pressure in the legs and uses position sensors which keep the table top at a fixed position relative to the floor. With this environmental isolation, the cavity does not appear to degrade the spectral resolution of the Ca spectroscopy at the 100 Hz level, but the cavity shows laser frequency drifts of $>$ 1 kHz/min. Furthermore, these drifts change direction on a 12 hour time scale, indicating less than ideal temperature stabilization of the cavity housing. The present drift is annoying but tolerable, because we can compensate for it. Since the small thermal mass of the 5 cm long cavity enhances its sensitivity to temperature variations, we will soon install a more massive cavity, which should reduce these drifts.

We tune our laser frequency with a double-passed AOM located between the beamsplitter and the cavity. By changing the frequency of the synthesizer which controls the AOM's frequency shift, we effectively tune the optical frequency of the laser light relative to the cavity fringe with synthesizer stability. Because the Fabry-Perot cavity resonance lying closest to the Ca resonance frequency is 1.3 GHz away, we use a 650 MHz modulator, which has a 200 MHz tuning range. We are thus able to scan the laser frequency over the 657 nm resonance in a repeatable and straightforward way simply by tuning the synthesizer frequency. We also are able to cancel residual cavity drifts to a large degree by including a slow ramp to the triangular wave used to drive the synthesizer for sweeping the laser frequency over the Ca resonance. Alternatively, we can suppress the drift by locking to a separate Ca atomic beam spectrometer (a similar approach was used for high resolution Mg spectroscopy [1]).

The total output power of the 657 nm ECDL is 4.5 mW, which is less than desired for the Ca spectroscopy. To increase the output power, we send 3 mW of this light through an optical amplifier which provides > 200 mW at 657 nm. Unfortunately, only ∼ 50% of this power is in the delta function of the laser frequency, while the rest is spread over 20 nm as amplified spontaneous emission. Moreover, the spatial mode of the amplifier contains much structure, so after spatial filtering with a polarization-preserving optical fiber, we have only 35 mW of useful output. While this power efficiency is less than ideal, 35 mW is still a significant improvement over the ECDL power, and the amplifier is extremely easy to use. We are also able to use the amplifier as an intensity stabilizer by controlling its injection current to keep the optical power emerging from the optical fiber constant. We have verified that the phase noise written onto the light due to the current modulation of the amplifier is less than 0.03 rad on our spectroscopic time scale. Thus it negligibly degrades the laser coherence for our measurements. In summary, our ECDL-amplifier system provides 30 mW of intensity-stabilized ($\Delta P/P \sim 10^{-3}$), frequency-stabilized ($\Delta\nu < 50$ Hz), and spatially filtered light, ideal for the high resolution spectroscopic investigations which will be described in subsequent sections.

2.5 Summary of apparatus design

A few comments summarizing the whole apparatus are in order here, because we opened the paper touting the potential compactness and robustness of the diode laser-based system. Without much effort to keep our optical set-up compact, we have used only 2/3 of a 1.5 m × 3.6 m optical table. This table also includes a Ca atomic beam apparatus with room left over for visiting transportable Ca frequency references and fiber coupling to other experiments. There is no reason why the size of the trapped-atom frequency reference could not be reduced by a factor of 2 or more. In terms of robustness, this system has run for more than a year and a half with little down time and can usually provide reliable Ramsey fringes within 30 minutes after turn-on.

3 Spectroscopic approach – Optical Ramsey fringes with shelving detection

3.1 High resolution spectroscopic technique – Time-domain pulsed optical Ramsey fringes

In order to probe the cold atoms in the absence of perturbing laser beams and magnetic gradient fields, it is necessary to use a pulsed measurement sequence. A typical sequence consists of three periods: trapping, excitation, and detection. Naturally, optimization of this measurement sequence is an essential part of our work, so we have investigated several excitation and detection strategies, which we will now discuss in some detail.

To excite the weak 657 nm transition we use the spectroscopic technique of optical Ramsey fringes, which is the optical analog of the microwave technique of separated oscillatory fields developed by Ramsey for use with atomic beams [32]. To appreciate the power of this technique, let us first consider some simpler excitation strategies and see where they fall short. If we excite the atoms with a single red pulse and detect the induced fluorescence, we are sensitive to the first-order Doppler shift. Due to the several milli-kelvin temperature of the atoms, this leads to linewidths of several megahertz (recall Fig. 3). In order to attain kilohertz linewidths, we must instead employ a sub-Doppler technique, the simplest of which consists of two π pulses, one from each direction. Such a technique produces a "saturation" dip in the center of the velocity distribution, since the atoms around zero velocity will be resonant with both pulses. The width of this dip is simply determined by Fourier transform of a single pulse, whose spectrum determines the range of velocities accessed by the excitation light for a given detuning. Then, by making our pulses longer, we can achieve narrower dips and in principle attain the desired kilohertz level linewidths. The drawback to this approach is that as we go to narrower linewidths, we excite a smaller range of velocities, and thus fewer atoms contribute to the saturation dip signal. For our distribution of velocities, a 1 kilohertz-wide signal would use only 1 atom in 1000, seriously degrading the signal-to-noise ratio.

An elegant solution to this problem is provided by the optical version of Ramsey's technique, which, as Bordé pointed out, is a type of atomic interferometer [33]. This method uses four $\pi/2$ pulses, two in one direction followed by two in the opposite. The need for four pulses rather than two (as in the microwave case) is a result of the much shorter wavelength of the optical radiation [15–17]. Let us now consider the excitation induced by this four pulse sequence. The Fourier transform of a pair of square pulses of duration τ and separated by time T is simply a $\cos(2\pi\Delta T)$ with a $\mathrm{sinc}^2(2\pi\Delta\tau)$ envelope centered around the laser frequency, where Δ is the detuning relative to the laser frequency. This spectral pattern is completely analogous to the Young double-slit interference pattern. One way to interpret this excitation process is that the first pair of pulses excites a velocity "grating" with a period $1/T$ in the atomic velocity distribution. In fact, two gratings are written in the velocity distribution, one in the ground state population, and one in the excited state population. Relative to the ground-state grating, the excited-state grating is actually shifted in optical frequency space by 23.1 kHz due to the absorption of a photon. To complete the excitation, a second pair of pulses (also separated by time T) then comes from the opposite direction and effectively reads out these gratings. The laser detuning will then determine how the read-out grating overlaps with a given population grating, yielding either a maximum or a minimum or something between. The resultant excited-state population is a sinusoidal function of the frequency with period $1/(2T)$ (and FWHM linewidth of $1/(4T)$) and a minimum located at line center. In reality, we have

two sets of fringes split by 23.1 kHz, one associated with the ground-state grating and one with the excited-state grating. At high resolution it is critical to choose Ramsey times T which allow these two patterns to add constructively (fringe periods which divide evenly into 23.1 kHz) in order to maintain maximum fringe contrast.

The beauty of this technique is that while the resolution is determined by the time between two pulses, the range of velocity classes participating is determined by the Fourier transform of a single pulse. Thus short pulses use a wide range of atomic velocity classes even when going to high resolution, and a high signal-to-noise ratio is maintained. In our case, a $\pi/2$ pulse needs to be $\sim 0.2\,\mu$s long to cover the whole velocity distribution (recall that the FWHM of our Doppler distribution is $\sim$ 3 MHz). Unfortunately, we have only enough red power to make a $\pi/2$ pulse in $0.5\,\mu$s, so at present we use only about 40% of our atoms. While this four-pulse sequence is a good choice for exciting the atoms, there may be still better choices which use more pulses and/or frequencies. Ertmer *et al.* have obtained Fabry-Perot-like resonances using multi-pulse sequences, which allow the use of shorter pulses while maintaining good contrast [34].

To generate the red excitation pulses we send the 657 nm light through two AOMs which enable us to send our $0.5\,\mu$s pulses in either of two counter-propagating directions (see Fig. 6). With such short pulses our red light is actually used for a total of only $2\,\mu$s out of the 5 ms measurement cycle, or less than 0.1% of the time. Therefore we have configured our optics so that when the AOMs are turned off, the light goes to our Ca atomic beam apparatus, where we generate Ramsey fringes for comparison with the trap fringes. See reference [1] for a similar arrangement.

3.2 Measurement cycle and data acquisition

Now that we have described the excitation technique, let us understand the basic measurement sequence before considering the more advanced version which incorporates shelving detection. The sequence commences with a 5−10 ms atom collection period, during which the trap fills with $\sim 10^7$ atoms. After the blue trap fields are extinguished, we excite the atoms with our sequence of four 657 nm pulses, using linearly polarized (parallel to the magnetic field) light to excite the $m = 0 \rightarrow m = 0$ transition. The measurement sequence concludes with a detection period during which we collect 657 nm fluorescence to measure the excitation. This fluorescence induces a PMT current which we measure with a gated integrator. We then repeat this measurement cycle as many times as desired before stepping the red frequency to a new frequency for the next data point. We control the length and functions of each step in the measurement cycle with a digital pulse generator. The pulse generator allows us to have as many as 1024 steps per cycle (we typically use 10–12) with 16 different digital control channels. With this device we can easily change or add steps, which greatly simplifies optimization of the system.

We acquire our data two different ways. During most of the early studies we used a computer-controlled system, which read the integrated PMT current at the end of a preset number of cycles (typically 100) for a given frequency. More recently, we have been taking data by slowly sweeping the red frequency with a synthesizer (swept by a function generator at $\sim$ 1 Hz) while running the measurement cycle continuously. We then acquire the data with a combination of a gated integrator and a sample-and-hold circuit. The resultant signal is stored (and averaged) on a digital oscilloscope. This differs from the first method in that we take sweeps much faster, although with less averaging per point. We can then average as many sweeps as desired on the digital oscilloscope to produce our final signal. In principle the two data collection techniques should give equivalent results, but we actually see less low frequency noise using the second technique due in part to atomic number fluctuations in the trap. The oscilloscope technique also enables one to monitor the "live" signal, which vastly expedites the optimization of various system parameters.

3.3 Shelving detection scheme

We have experimented with two strategies for detecting the excitation induced by the red pulses. The first, described above, was simple detection of the red fluorescence during the $400\,\mu$s decay time of the excited $^3P_1$ state. While this approach produced visible Ramsey fringes, the signal-to-noise ratio was low for two main reasons. The first was due to our poor solid angle coverage, which when coupled with the small quantum efficiency of photomultipliers at 657 nm ($< 5\%$) led to detection of only a few decays in a thousand. Second, the signal was sensitive to fluctuations in the number of trapped atoms from shot to shot, which are more than an order of magnitude above shot noise.

To address these issues, we now use a technique first applied to trapped ion spectroscopy [2], where a strong transition is used to monitor level populations which are modified by a weak transition. In our application of this shelving technique, we use a nearly resonant 423 nm standing wave pulse to measure the ground state population before and after red excitation. These blue pulses are $20\,\mu$s long and propagate collinearly with the red probes. We include these pulses in a modified shelving measurement sequence. This sequence uses the first blue probe pulse directly after trap turn-off (see Fig. 7). During this pulse we collect the induced blue fluorescence, which is proportional to ground state population (after background subtraction, which is $\sim 10\%$ of the signal). Next follows the red excitation sequence which transfers some fraction (typically $\sim 25\%$) of the population to the $^3P_1$ excited state, from which it will decay with a $400\,\mu$s time constant. However, during the first $20\,\mu$s of this decay time we illuminate the atoms with a second blue pulse to measure the population remaining in the ground state. After background subtraction, the ratio of fluorescence collected during the second blue pulse to that during the first yields the fraction

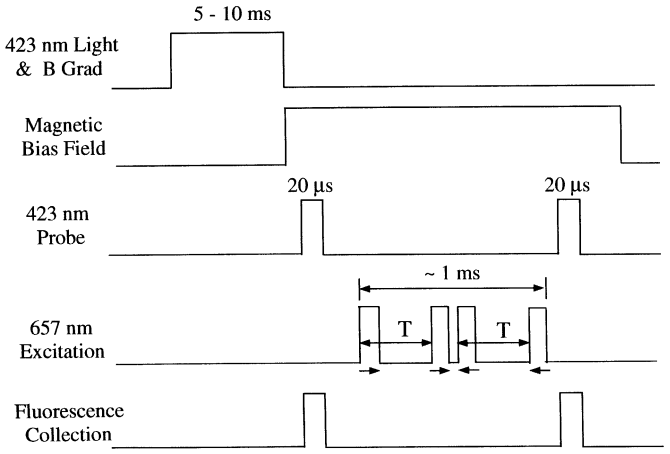

Fig. 7. Measurement cycle for generation of time-domain optical Ramsey fringes with shelving detection.

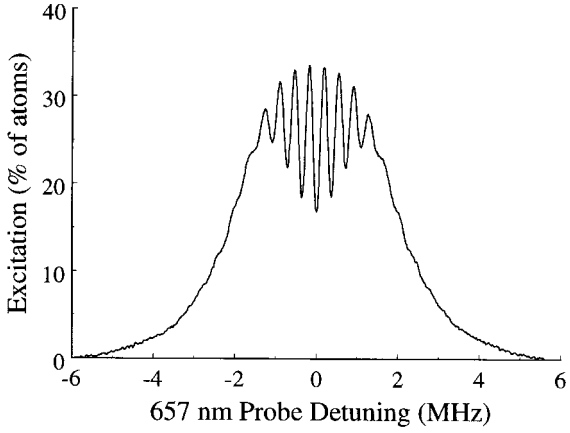

Fig. 8. Low-resolution optical Ramsey fringes taken with 0.5 μs pulse length, 1 μs pulse separation, and a probe intensity of 320 mW/cm^2. Trace is a single 2 s scan.

of atoms excited by the red pulse sequence. We perform this division either in software (for computer-controlled scans) or hardware (for live scans), and the resultant excitation is then expressed in terms of the fraction of atoms excited.

Shelving improves our signal-to-noise ratio in two ways. First, we can cycle tens or hundreds of photons per atom with the strong blue transition, which coupled with the higher PMT efficiency at 423 nm (23%) brings us close to the limit in which we detect one photon per atom per measurement cycle. In this "atom projection noise" regime [35], the quantum noise is determined not by the shot noise of the collected photons, but by that of the atom number itself. This limit offers us potentially a factor of > 20 improvement over the red detection shot noise limit. Second, our ratio or normalization procedure reduces the sensitivity of our blue detection system to trap number fluctuations, blue intensity and frequency noise, and other detection noise at frequencies below roughly 1 kHz (1 ms is typically the longest time between blue pulses, while we usually work at a few hundred microseconds). Similar normalization schemes have been used in other trap-based measurements to improve the signal-to-noise ratio [3,36]. We have to be a little more careful in our particular implementation of this technique, because we do not want to heat the atoms significantly with the first blue pulse. For this reason we use carefully balanced counter-propagating blue probes to form a true standing wave. Moreover, we tune the frequency of the probe slightly to the red of the transition ($\sim$ 15 MHz) and keep its total intensity down to 1.1 mW/cm^2. The resulting lower excitation rate prevents us from reaching the atom projection noise limit by a factor of 5, but it allows us to keep our atom temperatures within a factor of 2–3 of the Doppler limit. Because of this limitation, we have thus far seen an improvement in the signal-to-noise ratio of only 12 over red detection, but this should increase with the installation of our new fluorescence collection system, which should enable us to reach the atom projection noise limit.

4 Spectroscopic results

4.1 General features and recoil suppression

As a first example of the optical Ramsey fringes derived from trapped atoms, in Figure 8 we show a scan of the whole velocity distribution taken at very low Ramsey resolution. Because we took this data with the shelving technique, the excitation is expressed in terms of the fraction of the atoms excited. The width of the background is 4 MHz, representing a convolution of the Fourier transform of a single 0.5 μs red pulse and the 2 mK velocity distribution. The fringes are superimposed on a saturation dip of width 2 MHz, also determined by Fourier transform of a single red pulse. At this low resolution (linewidth = 165 kHz) the recoil is not resolved, so both recoil components add constructively. With the improvement due to the implementation of our shelving detection scheme, we can have good a signal-to-noise ratio with a single 2 s trace. Also, the Ramsey wiggles persist even into the wings of the velocity distribution. We think that these wiggles (which persist in computer simulations as well) are the result of excitation by the sidelobes of the $\mathrm{sinc}^2(2\pi\Delta\tau)$ spectral envelope of our excitation pulses.

To go to higher resolution we simply increase the Ramsey time separation T for each pair of pulses. Figure 9 shows a trace taken with a Ramsey time of 260 μs, which corresponds to a resolution of 960 Hz. Here, we have chosen a resolution for which the two recoil components add constructively to yield the greatest contrast. Alternatively, one can suppress the higher frequency set of fringes by adding a blue pulse in between red pulses 2 and 3. This pulse effectively washes out the velocity grating written in the ground state by cycling photons between the ground state and the $^1\mathrm{P}_1$ state. With the grating absent, the subsequent pair of red pulses cannot produce the fringe structure for this recoil component. Such "recoil suppression" has been demonstrated by several groups [37–39] and is easily implemented here by using an extra pulse from our blue standing wave probe system. We can then use

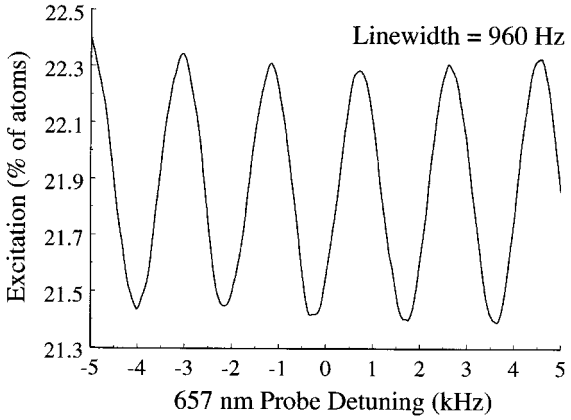

Fig. 9. Sub-kilohertz optical Ramsey fringes taken with 0.5 μs pulse length and a 260 μs pulse separation. This fringe pattern is the average of 36, 1 s sweeps.

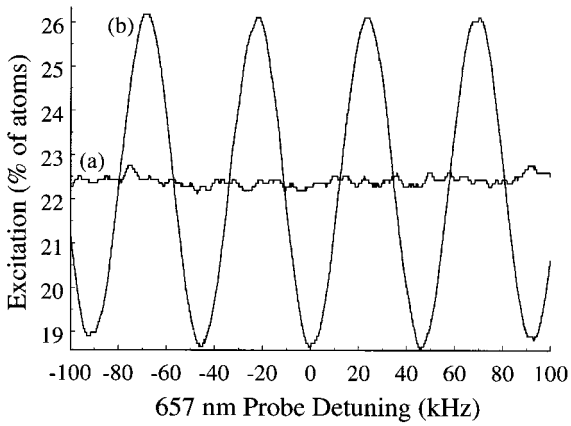

Fig. 10. Recoil suppression (curve (b)) through use of a single 10 μs blue probe pulse between red pulses 2 and 3, which reveals the fringes due to a single recoil component. Curve (a) was taken without recoil suppression and shows an absence of fringe structure due to destructive superposition of the two recoil components.

any desired resolution without loss of contrast, although with a signal only half as large as when both components add constructively. We see a nice demonstration of recoil suppression in Figure 10, which shows frequency scans taken at a resolution of 23.1 kHz under slightly different conditions. Sweep (a) was taken with usual excitation sequence, but at this resolution the two Ramsey patterns cancel and no fringes are visible. For sweep (b) we added a 10 μs blue standing wave pulse, which effectively suppressed the high frequency contribution, thus revealing the fringes associated with the low frequency recoil component. In practice we rarely use this technique, because the restriction on our choice of resolution does not pose a great limitation, and we prefer the larger signal resulting from the inclusion of both recoil components. However, the capability of suppressing a recoil component will be useful when evaluating systematic shifts in the frequency standard.

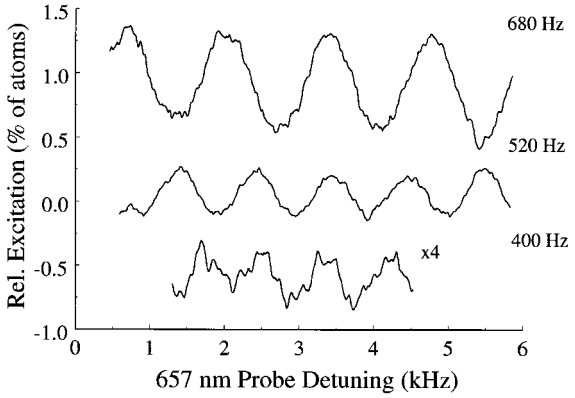

Fig. 11. Ultra-high resolution fringes at three different resolutions including 400 Hz, the natural linewidth. Note the decay of the fringe amplitude at higher resolution. These fringe patterns are each the average of 50, 1 s sweeps.

Closer examination of the 960 Hz linewidth fringes in Figure 9 reveals several important features. We see that our apparatus can achieve an excellent signal-to-noise ratio at sub-kilohertz resolution with only 30 s of averaging time. This signal-to-noise ratio implies a fractional frequency stability of $\sim 10^{-14}$ in 1 s. However, there is also evidence of excess low-frequency noise (the amplitude of the fringes should be equal), which persists despite the normalization process. This residual noise, whose spectrum peaks between 15 and 30 Hz, is due in part to acoustic noise and vibrations present in the lab, and will have to be reduced to achieve shot-noise-limited performance for the frequency standard.

4.2 High resolution spectroscopy

In Figure 11 we see Ramsey fringes taken at several higher resolutions. Shown here are just the central three fringes (out of $\sim$ 1000 in the pattern). We see that we can obtain visible fringes well below the kilohertz limit. When we improve our signal-to-noise ratio further, we can even imagine going to resolutions below the 400 Hz natural linewidth, although with greatly reduced contrast due to the effects of spontaneous emission. In fact, there is already a marked loss in fringe amplitude with increase in resolution (see Fig. 12). In reducing the linewidth from 10 kHz to 520 Hz, we see a factor of 20 reduction in fringe size, of which only a factor of 3 can be attributed to spontaneous decay from the excited state.

There are in fact several technical problems which contribute to the loss of signal. The biggest culprit is the residual atomic temperature, which leads to a ballistic expansion of the atom cloud of several millimeters during the probe period used for the highest resolutions. To have good wavefronts and equal intensities for all atoms thus requires the use of large probe beams. The use of large probe beams, however, reduces our already insufficient 657 nm probe intensity. We have found a compromise by using 2 mm diameter probe beams which give good

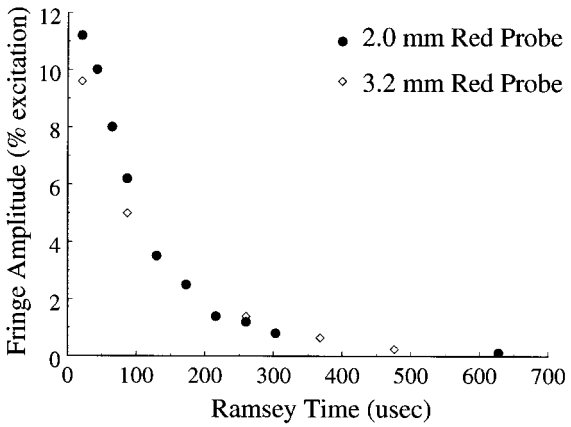

Fig. 12. Decay of fringe amplitude with Ramsey time T for two different beam diameters.

signal-to-noise ratio, although this may lead to increased systematic effects due to the larger range of k-vectors sampled by the atoms. This smaller size does lead to a more rapid rate of fringe amplitude loss, but the amplitude is still larger at low resolution and is comparable at high resolution.

Another potential contributor to the loss of fringe amplitude is the lack of sufficient coherence for the 657 nm probe light. While the fast laser linewidth appears to be sufficiently narrow for these measurements, we have not yet been able to rule out decoherence due to cavity fluctuations. Implementing an improved reference cavity should reduce these contributions and enable better measurements of the fluctuation spectrum.

As a first test of the frequency stability of our Ca system, we have locked our 657 nm laser to the central Ramsey fringe by feeding back to the double-passed AOM between the laser and the Fabry-Perot cavity. (To identify the central fringe we can take scans at several different resolutions and observe where the fringe minima line-up [12].) With a modulation frequency of 100 Hz and a (square-wave) modulation width equal to the 3 kHz fringe linewidth chosen for this measurement, we measure an Allan variance of 5×10^{-14} in 1 s with a servo bandwidth of $\sim$ 10 Hz. The frequency fluctuations were measured relative to our reference cavity with the cavity drift suppressed, and so represents a useful but indirect measurement of the stability provided by the lock to Ca. The Allan variance is a factor of 15 larger than that corresponding to the atom projection noise limit for our current fringe amplitude. Of course we do not expect to reach atom projection noise since we have not yet achieved the 1 photon per atom per measurement cycle regime, and we still have some residual low-frequency noise.

5 Future improvements and applications

We have described here an all-diode-based Ca trap and spectroscopy system, which could serve as the starting point for numerous experimental investigations including precision atomic interferometric measurements and cold collision studies. Our work focuses on the application of Ca to optical frequency standards. With this, our initial effort, we have achieved a fractional frequency stability competitive with other state-of-the-art systems, and there are certainly good prospects for much improvement. One aspect of our future work will concentrate on improving the short term performance of our apparatus. Most of the present noise contributors are readily addressed and will soon be reduced. Imminent improvements include increased solid angle coverage for the fluorescence detection system, optimization of the normalization, and an improved Fabry-Perot reference cavity for the 657 nm locking system. We also plan to incorporate the 672 nm repumping light into the trap system to increase the number of atoms in the trap. Another critical aspect of the development of this frequency reference will be the minimization of systematic errors. We have reduced the contributions due to stray light to the sub-hertz level, but we still have to implement our new magnetic field gradient chopper to reduce the second order Zeeman shifts.

These improvements should help us to approach the remarkable potential of this transition. With 10^8 atoms, a contrast amplitude of 20% (in terms of the total atom number), and a 10 ms cycle time, the shot-noise-limited stability at 1 kHz resolution would be an impressive 10^{-16} in 1 s. To reach this fringe amplitude (a factor of 10 higher than we have now at 1 kHz), however, will require reduced velocity of the atomic sample and/or higher optical power. The rms velocity of our atomic sample is currently 1.5–2 times larger than the Doppler limit but has so far defied our attempts at further reduction. Moreover, to achieve the full potential of this transition, colder temperatures than this broad blue cooling transition can provide may well be required. Cooling on the narrow 657 nm transition could potentially lead to much colder temperatures, but is too inefficient due to the long lifetime of the excited state (unlike the analogous transition in Sr, which has shown spectacular cooling performance [40,41]). Other cooling schemes such as two-photon cooling on the $(4s^2)^1\mathrm{S}_0$–$(4s5s)^1\mathrm{S}_0$ transition or using a closed transition between higher lying triplet states may provide the best alternative.

Another important direction this work will take will be the connection of the Ca optical frequency to other frequency domains through the use of phase-locked frequency chains. As a first application, we are presently connecting the Ca 657 nm transition to the 282 nm optical frequency standard based on a trapped Hg^+ ion [8]. Using the value measured by PTB for Ca [19] will enable us to give a good frequency value for the Hg^+ transition. We plan ultimately to connect both of these systems by a phase-coherent frequency chain to the Cs microwave standard at 9 GHz.

The authors gratefully acknowledge valuable experimental assistance from Michelle Stephens and Carl Weimer in the early stages of the experiment. We also thank J.C. Bergquist, H. Robinson, J.L. Hall, K.R. Vogel, G. Zinner, T. Trebst, P. Kersten, F. Riehle, and J. Helmcke for their contributions

and helpful comments on this research, and T. Zibrova for the excellent coatings for the windows and beamsplitters. We are all indebted to the contributions to this field by the late Richard Barger, whose work guided the way and whose input during the early phases of our project was invaluable. Finally, we acknowledge the support of the Air Force Office of Scientific Research and the National Research Council. *Contribution of the U.S. Government; not subject to copyright.*

References

1. K. Sengstock, U. Sterr, J.H. Müller, V. Rieger, D. Bettermann, W. Ertmer, Appl. Phys. B **59**, 99 (1994); D.J. Wineland, J.C. Bergquist, W.M. Itano, R.E. Drullinger, Opt. Lett. **5**, 245 (1980).
2. For a different approach to electron shelving detection on this transition, see T. Kurosu, G. Zinner, T Trebst, F. Riehle, Phys. Rev. A **58**, R4275 (1998).
3. See for example, J.L. Hall, M. Zhu, P. Buch, J. Opt. Soc. Am. B **6**, 2194 (1989).
4. S. Ghezali, Ph. Laurent, S.N. Lea, A. Clairon, Europhys. Lett. **36**, 25 (1996).
5. Th. Kisters, K. Zeiske, F. Riehle, J. Helmcke, Appl. Phys. B **59**, 89 (1994).
6. Jun Ye, Long-Sheng Ma, J.L. Hall, IEEE Trans. Instr. Meas. **46**, 178 (1997).
7. Th. Udem, A. Huber, B. Gross, J. Reichert, M. Prevedelli, M. Weitz, T.W. Hänsch, Phys. Rev. Lett. **79**, 2646 (1997).
8. J.C. Bergquist, W.M. Itano, F. Elsner, M.G. Raizen, D.J. Wineland, in *Light Induced Kinetic Effects on Atoms, Ions, and Molecules*, edited by L. Moi, S. Gozzini, C. Gabbanini, E. Arimondo, F. Strumia (ETS Editrice, Pisa, 1991), p. 291.
9. A.A. Madej, J.E. Bernard, B.G. Whitford, L. Marmet, K.J. Siemsen, in *1998 Conference on Precision Electromagnetic Measurements Conference Digest*, July 6-10, edited by T.L. Nelson, p. 323.
10. A.M. Akulshin, A.A. Celikov, V.L. Velichansky, Opt. Commun. **93**, 54 (1992).
11. G.M. Tino, M. Barsanti, M. de Angelis, L. Gianfrani, M. Ingusio, Appl. Phys. B **55**, 397 (1992); A.A. Celikov, A.M. Akulshin, V.L. Velichansky, A.S. Zibrov, Laser Phys. **5**, 739 (1995).
12. G. Zinner, Ph.D. thesis, University of Hannover, PTB-Bericht Opt-58, Braunschweig, p. 54, 1998.
13. F. Ruschewitz, J.L. Peng, H. Hinderthur, N. Schaffrath, K. Sengstock, W. Ertmer, Phys. Rev. Lett. **80**, 3173 (1998).
14. C.W. Oates, F. Bondu, L. Hollberg, in *Sixteenth International Conference on Atomic Physics*, August 3-7, 1998, edited by W.E. Baylis, G.W.F. Drake, p. 115.
15. R.L. Barger, J.C. Bergquist, T.C. English, D.J. Glaze, Appl. Phys. Lett. **34**, 12 (1979).
16. Ye.V. Baklanov, B.Ya. Dubetsky, V.P. Chebotayev, Appl. Phys. **9**, 171 (1976).
17. J.C. Bergquist, S.A. Lee, J.L. Hall, Phys. Rev. Lett. **38**, 159 (1977).
18. Ch.J. Bordé, Ch. Salomon, S. Avrillier, A. van Lerberghe, Ch. Bréant, D. Bassi, G. Scoles, Phys. Rev. A **30**, 1836 (1984).
19. H. Schnatz, B. Lipphardt, J. Helmcke, F. Riehle, G. Zinner, Phys. Rev. Lett. **76**, 12 (1996).
20. T. Quinn, Metrologia **30**, 523 (1994).
21. C. Monroe, W. Swann, H. Robinson, C. Wieman, Phys. Rev. Lett. **65**, 1571, (1990).
22. We note with admiration the vapor cell trap constructed for Sr – see T.P. Dinneen, K.R. Vogel, E. Arimondo, J.L. Hall, A. Gallagher, Phys. Rev. A **59**, 1216 (1999).
23. Takayuki Kurosu, Fujio Shimizu, Jpn J. Appl. Phys. **29**, L2127 (1990).
24. B.P. Anderson, M.A. Kasevich, Phys. Rev. A **50**, R3581 (1994).
25. C.W. Oates, M. Stephens, L. Hollberg, in *Proc. of the 1997 IEEE Freq. Cont. Symp.*, 28-30 May, 1997, edited by J. Vig, p. 219.
26. A. Noble, M. Kasevich, Rev. Sci. Inst. **65**, 3042 (1994).
27. K. Lindquist, M. Stephens, C. Wieman, Phys. Rev. A **46**, 4082 (1992).
28. We initially used a thin wire placed 5 cm downstream from the oven as a shield for the trapped atoms. This proved very useful in our first observations of the trap as it provided a dark stripe in the background fluorescence against which to view the (not yet optimized) trapped sample. Subsequent removal of the wire doubled the loading rate.
29. Mention of specific products is for technical clarity and does not imply endorsement by NIST. Other products may be even more suitable.
30. K. Gibble, S. Chu, Phys. Rev. Lett. **70**, 1771 (1993).
31. R.W.P. Drever, J.L. Hall, F.V. Kowalski, J. Hough, G.M. Ford, A.J. Munley, H. Ward, Appl. Phys. B **31**, 97 (1983).
32. N.F. Ramsey, Phys. Rev. **78**, 695 (1950).
33. C.J. Bordé, Phys. Lett. A **140**, 10 (1989); for beautiful demonstrations of time-domain optical Ramsey spectroscopy with trapped atoms, please see references [1,5].
34. H. Hinderthür, A. Pautz, V. Rieger, F. Ruschewitz, J.L. Peng, K. Sengstock, W. Ertmer, Phys. Rev. A **56**, 2085 (1997).
35. D.J. Wineland, W.M. Itano, J.C. Bergquist, F.L. Walls, in *Proceedings of the 35th Ann. Symp. Freq. Control*, May 1981, p. 602.
36. C.W. Oates, K.R. Vogel, J.L. Hall, Phys. Rev. Lett. **76**, 2866 (1995).
37. F. Riehle, J. Ishikawa, J. Helmcke, Phys. Rev. Lett. **61**, 2092 (1998).
38. T. Kurosu, A. Morinaga, Phys. Rev. A **45**, 4799 (1992).
39. U. Sterr, K. Sengstock, J.H. Müller, D. Bettermann, W. Ertmer, Appl. Phys. B **54**, 341 (1992).
40. K.R. Vogel, T.P. Dinneen, A. Gallagher, J.L. Hall, in *1998 Conference on Precision Electromagnetic Measurements Conference Digest*, July 6-10, edited by T.L. Nelson, p. 303.
41. H. Katori, T. Ido, Y. Isoya, M. Damion, M. Kuwata-Gonodami, in *Sixteenth International Conference on Atomic Physics*, August 3-7, 1998, edited by W.E. Baylis, G.W.F. Drake, p. 104.