An Optical Lattice Clock with Neutral Strontium

Richard Hobson
Balliol College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

2016

Acknowledgements

I'm immensely grateful to my colleagues on the Sr lattice clock team – Ian Hill and William Bowden. Without their countless hours sacrificed to the lab, our experiment could never have become the slick metrology machine it is today. Many thanks also go to all the others who have made significant contributions to the project over the years: Sean Donnellan, Liz Bridge, Anne Curtis, Yuri Ovchinnikov, Marco Menchetti, Matt Earnshaw, Joshua Hughes, and Daryl Paul.

Beyond the lattice clock lab, many more people have contributed essential measurement infrastructure within the NPL Time and Frequency group. In particular, I thank Steve King, Peter Nisbet-Jones, Rachel Godun, Jonathan Jones, Charles Baynham, Helen Margolis, Antoine Rolland, Fred Baynes, and Peter Whibberley, both for their work and for friendly advice and discussions.

Of course, none of this research would have been possible without the financial support of the EPSRC and the National Metrology Service. If you continue to fund lattice clocks, all of your wildest dreams will come true.

Finally, I credit my supervisors – Patrick Gill at NPL and Patrick Baird at Oxford – for their consistent support of the project and my studentship behind the scenes.

Abstract

Strontium optical lattice clocks can measure the passage of time with extraordinary precision. Capitalising on this precision, we can anticipate the deployment of lattice clocks to search for new physics beyond the standard model, to build new technologies for geodesy and navigation, and potentially to underpin a future redefinition of the SI second.

This thesis reports on the construction and evaluation of a robust Sr optical lattice clock at NPL. We describe the apparatus needed to capture, cool, and load samples of neutral strontium atoms into a magic-wavelength, far off-resonant lattice trap at 813 nm. We provide details of our optical local oscillator – the "clock laser" – and how it is used to realise an Sr-referenced optical frequency standard. We rigorously characterise the various contributing factors which limit the performance of the clock, focusing on the standard measures of (1) frequency instability, and (2) systematic frequency uncertainty. Finally, we introduce new innovations for improving the accuracy of the ^{88}Sr lattice clock, including methods of "modified hyper-Ramsey" and multi-photon spectroscopy of the clock transition.

Contents

3 Clock stability 65

List of Figures

Chapter 1

Introduction

Atomic clocks are incredibly precise timing machines. In particular, the best strontium-atom optical lattice clocks can measure the passing of time to the 18^{th} decimal place, implying that it would take them 15 billion years to accumulate just 1 second of error [193]. Alongside atomic magnetometers, accelerometers, and gyroscopes, these exceptionally accurate clocks belong to a class of emerging quantum technologies that are now beginning to reach maturity after the advent of laser cooling of atoms in the 1970s [288]. These technologies give us the ability to sense and characterize the world around us with unparalleled precision.

1.1 What is clocks?

Most clocks share a similar overall structure comprising: (1) a local oscillator, (2) a counter, and (3) a frequency reference.

The role of the local oscillator is to provide the 'ticks' of the clock by evolving in regular, deterministic cycle. Physically, it could be anything from a swinging pendulum in a grandfather clock, an electronic oscillator in a watch or phone, or a maser or laser in an atomic clock, but the important thing is that each cycle should mark the passing of a well-known chunk of time t_{cyc}. The purpose of the clock's counter is then to keep track of a measured 'clock time' $t_{\text{clock}} = N t_{\text{cyc}}$ by counting the number of cycles N that have occurred. Finally, to make sure that its tick rate doesn't drift, a good clock will also include a frequency reference, exemplified by the quartz crystal in an electronic watch, whose role is to posses a well-characterized resonance frequency ν_0 towards which we can constantly steer the local oscillator.

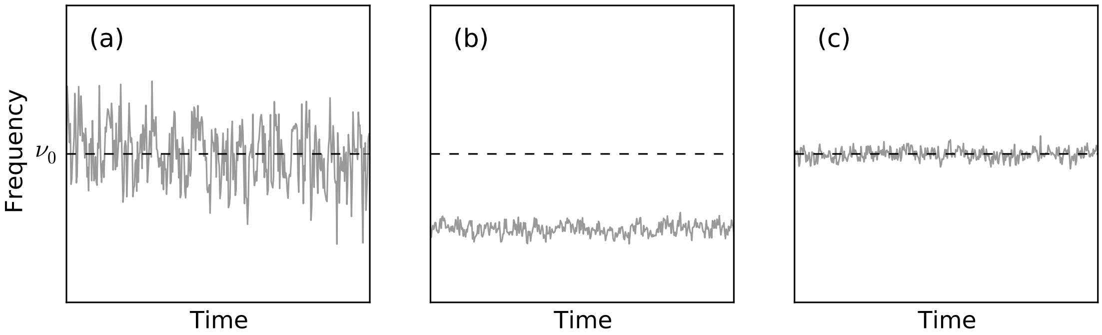

Figure 1.1: Distinction between accuracy and stability. Clock (a) is accurate but not stable, clock (b) is stable but not accurate, and clock (c) is both stable and accurate. For simplicity we have assumed the evaluated systematic offset is zero.

1.1.1 Atomic clocks

In an atomic clock, the frequency reference is either a single atom or an isolated sample of several atoms of the same species. The idea is to exploit a resonant frequency $\nu_0 = (E_2 - E_1)/\hbar$ corresponding to a particular 'clock transition' between two quantum states of the atom. By repeatedly probing the atoms using radiation from the clock's local oscillator (e.g. a laser or a microwave source), we can constantly steer the frequency ν of that local oscillator towards $\nu = \nu_0$ by trying to maximize the atom's resonant response. This ensures that the clock's tick rate is ultimately tied to an atomic resonance which, at least to a very good approximation, doesn't change over time.

Above all, this approach has one important advantage over quartz crystals or pendulums: reproducibility. Since all ^{133}Cs atoms (for example) are exactly the same regardless of their position in time or space, we can confidently expect two completely independent clocks to agree on their measure of time, without the need for mutual calibration or exchange of artifacts, just by using ^{133}Cs atoms as the reference in both clocks. Because of this, it was decided in 1967 actually to *define* time within the International System of Units (SI) using atomic clocks [90]:

> The second is the duration of 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the ^{133}Cs atom

However, try as we might to avoid them, there are always some external perturbations to the atom's resonant frequency from e.g. Zeeman, Stark, Doppler, and collisional effects, which cause real atomic clocks to tick slightly faster or slower than they should. Therefore, to be unambiguous, the definition of time must specify an 'ideal' atomic resonance ν_0 which is free from such perturbations – in the SI definition, for example, ν_0 corresponds to a completely isolated ^{133}Cs reference atom held stationary in a free-space radiation bath at a temperature of absolute zero.

Since the ideal conditions cannot be perfectly reproduced in real life, the perturbations of an atomic clock must be individually measured and characterized with the goal of estimating a total systematic frequency offset, ν_{sys}. This process of characterisation – the 'systematic evaluation' – allows us to compensate for the perturbations and restore the accuracy of the clock. However, even after careful characterisation there remains some residual uncertainty $\delta\nu_{\text{unc}}$ on our estimation, ultimately leading to timing errors which cannot be compensated for. Further, the steering of the local oscillator towards atom resonance is never perfect, so the clock will have a time-dependent frequency instability $\delta\nu(t)$ which can cause significant timing errors in the short term even if it averages to zero in the long term.

Summarizing all these effects, we can write the frequency of an atomic clock in the form:

$$\nu(t) = \nu_0 + \nu_{\text{sys}} + \delta\nu_{\text{error}} + \delta\nu(t) \tag{1.1}$$

where, assuming we haven't made any big mistakes when analysing all the systematic perturbations, $\delta\nu_{\text{error}}$ is a random variable with a mean of zero and a standard deviation equal to $\delta\nu_{\text{unc}}$.

Within this equation are the two most important metrics of performance for atomic clocks:

- The *relative uncertainty* $\delta\nu_{\text{unc}}$, sometimes called the relative (in)accuracy, which represents the timing error we'd expect the clock to build up in the long term relative to an ideal clock of the same atomic species.
- The *instability* $\delta\nu(t)$, which determines how long it takes to make a frequency measurement with a given precision – if high precision is required, then we have to wait a long time for the statistical fluctuations $\delta\nu(t)$ to average out.

A final metric of clock performance is the uncertainty in *absolute* frequency, or the 'absolute (in)accuracy'. This describes the uncertainty in the knowledge of the atomic clock frequency when expressed in SI units of Hz – i.e. the uncertainty of the clock relative to an ideal ^{133}Cs clock rather than relative to its own unperturbed resonance. In a ^{133}Cs clock, the systematic uncertainty and the absolute accuracy are one and the same thing, but in an atomic clock based on any other species (such as ^{87}Sr) the evaluation of absolute accuracy requires comparison against the ^{133}Cs primary standard.

1.2 Optical Atomic Clocks

1.2.1 A brief history

Optical transitions in trapped atoms were recognised as a possible basis for clocks with 10^{-18} fractional precision well before such systems became practical [66], even at a time when microwave atomic clocks were languishing at 10^{-13} levels of inaccuracy. The main reason for optimism was simple: since the tick rate ν_0 of an optical clock is 10^5 times faster than in a microwave clock, a given absolute frequency error $\delta\nu$ results in a much smaller fractional frequency error $\delta\nu/\nu_0$. However, the fast tick rate also presented a difficulty: where microwave signals can be straightforwardly counted electronically, oscillations above a THz are simply too fast to count directly.

An early solution to this problem was to build a phase-locked frequency multiplication chain, starting at a countable microwave frequency and multiplying up into the infrared/optical region for comparison against the optical clock under test. In each step of the chain a phase lock could be implemented by coupling both the n^{th} and $(n+1)^{\text{th}}$ oscillators onto a fast-response Schottky or metal-insulator-metal (MIM) diode, so that an RF beat could be detected between a harmonic of ν_n and the fundamental of ν_{n+1}. If the RF beats at every step are phase-locked to known RF frequencies, a well-defined frequency ratio can be inferred between the first (microwave) and the last (IR/optical) element in the chain.

Conceptually the frequency-chain approach was quite intuitive, but in practice a horrifyingly complex proliferation of bulky gas lasers and microwave electronics was needed in order to bridge the 4-5 orders of magnitude gap in frequency [123, 140]. As a result, only a small handful of optical frequency measurements were ever performed using such a chain, with typical uncertainties of a few parts in 10^{13} [24, 238].

Thankfully, the advent of broadband femtosecond frequency combs in 1999-2000 led to a much easier method of counting optical frequencies, entirely circumventing the need for a frequency chain. In early experiments, the caesium D1 [270] and hydrogen 1S-2S [196, 219] transitions could be accurately measured using a hybrid scheme of optical frequency division [263] complemented by a comb-based frequency bridge. Soon afterward, the first octave-spanning frequency comb [126] simplified the counting of optical frequencies even further − with the the carrier-envelope offset

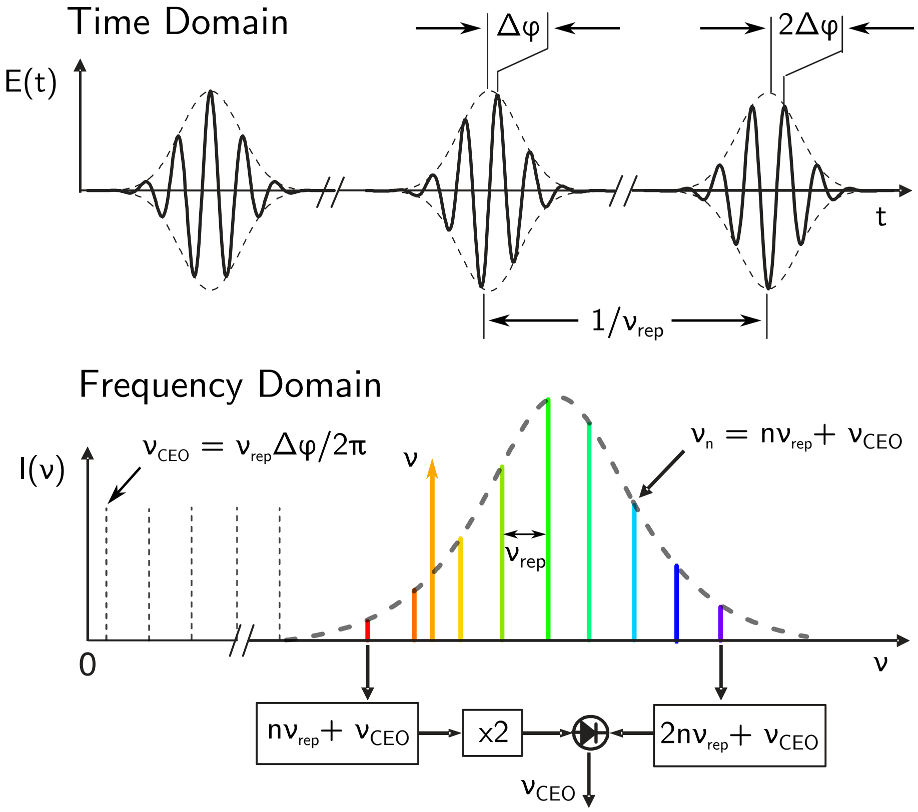

Figure 1.2: The self-referencing technique for an octave-spanning frequency comb. The comb output is a train of pulses in the time domain with repetition rate ν_{rep}. The carrier-envelope phase offset φ constantly evolves between pulses due to dispersion in the lasing cavity. If the comb spans a full octave in the frequency domain, the carrier-envelope offset evolution frequency ν_{CEO} can be extracted as shown in the figure. After stabilising the comb's ν_{CEO} and ν_{rep}, any optical frequency ν can be measured using the optical beat against the nearest comb mode. [Adapted from [112]]

ν_{CEO} fixed using the self-referencing technique depicted in figure 1.2, the comb could be used as a kind of 'ruler' to measure any optical frequency ν using the relation

$$\nu = n\nu_{\mathrm{rep}} + \nu_{\mathrm{CEO}} \pm \Delta \tag{1.2}$$

where ν_{rep} is the comb repetition rate and Δ is the beat signal between the optical frequency under measurement and the n^{th} comb mode.

The octave-spanning comb was first used for an absolute optical frequency measurement in 2000 [71], and further measurements quickly followed on Hg^+ [269] and Yb^+ quadrupole [248] clock transitions at an unprecedented 1×10^{-14} precision. Soon the new technology had unleashed the series of measurements plotted in figure 1.3

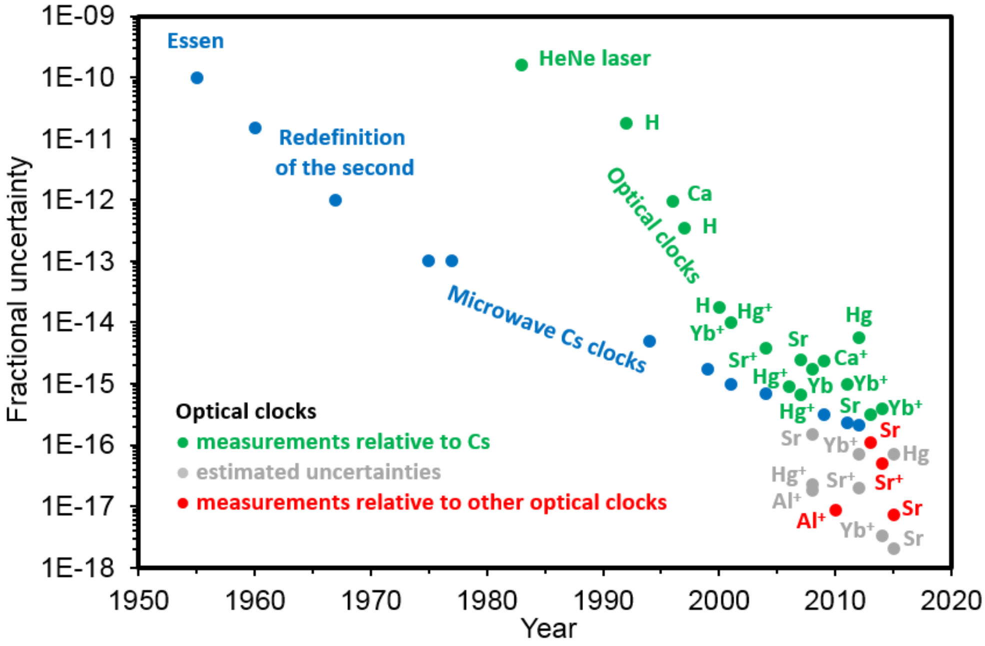

Figure 1.3: Historical progress in atomic clock accuracy. Courtesy of Rachel Godun.

[42, 57, 116, 137, 145, 152, 173, 178, 205, 211], with accuracy converging steadily towards the $\sim 2 \times 10^{-16}$ limit of the best Cs fountain standards [95, 104, 157]. Importantly, none of these measurements were limited in accuracy or stability by the combs themselves, and thus the true performance of the clocks under comparison was revealed – indeed, experiments show that a good pair of independent combs will typically agree to within 10^{-19} or even lower [125, 171].

The comb-fueled progress in optical clocks has now led to total systematic uncertainties in the 10^{-18} range for some optical clocks (see table 1.1), surpassing the best Cs fountains by two orders of magnitude. This is of course a good thing, but it also leads to a technical problem: since comparison against a Cs clock is currently by definition the only way to measure frequencies in SI units of Hz, it is impossible to measure the absolute (Hz) frequency of an optical clock with better accuracy than the best Cs fountain. This limitation has led to the investigation of direct frequency *ratios* between optical clocks, enabling much higher precision limited only by the performance of the worse of the two optical clocks. With several such ratios now possessing uncertainties below the so-called 'SI limit' of Cs-clock inaccuracy [160, 223, 294], a debate has begun about whether to formulate a new definition of the SI second (and thus the Hz) in terms of optical standards [89]. However, even though many optical clocks have been estimated by their makers to be very accurate

Table 1.1: Best reported accuracies and stabilities of optical clocks as of 2015

Species	Transition	λ (nm)	$\Delta\nu_{\mathrm{nat}}$ (Hz)	Uncertainty	Stability[a] $/\sqrt{\tau}$
$^{27}\mathrm{Al}^+$	$^1\mathrm{S}_0-{}^3\mathrm{P}_0$	268	0.008	8.6×10^{-18} [54]	2.0×10^{-15} [54]
$^{88}\mathrm{Sr}^+$	$^2\mathrm{S}_{\frac{1}{2}}-{}^2\mathrm{D}_{\frac{5}{2}}$	674	0.4	2×10^{-17} [172]	1.0×10^{-14} [172]
$^{199}\mathrm{Hg}^+$	$^2\mathrm{S}_{\frac{1}{2}}-{}^2\mathrm{D}_{\frac{5}{2}}$	282	1.7	1.9×10^{-17} [223]	2.8×10^{-15} [223]
$^{171}\mathrm{Yb}^+$	$^2\mathrm{S}_{\frac{1}{2}}-{}^2\mathrm{D}_{\frac{3}{2}}$	436	3.1	1.1×10^{-16} [114]	–
$^{171}\mathrm{Yb}^+$	$^2\mathrm{S}_{\frac{1}{2}}-{}^2\mathrm{F}_{\frac{7}{2}}$	467	10^{-9}	3.2×10^{-18} [118]	5.3×10^{-15} [114]
$^{87}\mathrm{Sr}$	$^1\mathrm{S}_0-{}^3\mathrm{P}_0$	698	0.001	2.2×10^{-18} [193]	1.6×10^{-16} [6]
$^{88}\mathrm{Sr}$	$^1\mathrm{S}_0-{}^3\mathrm{P}_0$	698	0	2.9×10^{-15} [4]	2.2×10^{-14} [4]
$^{171}\mathrm{Yb}$	$^1\mathrm{S}_0-{}^3\mathrm{P}_0$	578	0.01	3.4×10^{-16} [152]	3.2×10^{-16} [109]
$^{174}\mathrm{Yb}$	$^1\mathrm{S}_0-{}^3\mathrm{P}_0$	578	0	1.5×10^{-15} [211]	3×10^{-15} [211]
$^{199}\mathrm{Hg}$	$^1\mathrm{S}_0-{}^3\mathrm{P}_0$	266	0.1	7.2×10^{-17} [294]	3.0×10^{-15} [294]

---

[a] We describe clock stability using the Allan deviation $\sigma_y = X/\sqrt{\tau}$. See chapter 3 for definition.

(grey points of figure 1.3 [117, 118, 167, 172, 193, 223, 294]), only very few clocks have actually demonstrated reproducibility below 10^{-16} (red points [18, 54, 145, 160, 271]). More comparisons will most likely have to be made before a redefinition of the second is undertaken.

1.2.2 Who needs better clocks?

Caesium fountains are already accurate to an impressive 1 second in 150 million years. Why, then, are we trying to make even better atomic clocks? In this section I outline three arguments:

1. Atomic clocks provide ultra-precise tests of relativity, aiding the search for possible physics beyond the standard model;
2. Highly stable clocks could enable improved technologies in navigation, geodesy, and astronomy;
3. Atomic clocks are versatile platforms for investigating atomic physics – just how far can we go in the manipulation of the quantum states of cold atoms?

The combination of all these arguments, along with the fast pace of recent progress captured in figure 1.3, suggests that optical atomic clocks are likely to be a burgeoning field of research and technology for some years to come.

1.2.2.1 Fundamental physics

Local Position Invariance: Motivated by the predictions of some theories beyond the standard model, there is a long history of scientific inquiry into the possible variation of fundamental constants over time and across space [272]. If found, such variations would violate the principle within general relativity of Local Position Invariance (LPI), which states that the outcome of any local experiment should be independent of where and when it is performed. (Here, an experiment is 'local' if it occurs in a small enough region where the effect of curvature of space-time is negligible). A conclusive measurement of LPI violation, combined with details of its dependence on time and space, could help to refine the search for a Grand Unified Theory.

Atoms are singled out as sensitive tools for measuring LPI violation because their transition frequencies are determined by the local physics at the position of the atom. In particular, theoretical models of a range of transitions in different atomic species reveal different sensitivities to the fine structure constant α, the proton-to-electron mass ratio $\mu = m_p/m_e$, and the ratio m_q/Λ_{QCD} between the light quark mass and the quantum chromodynamics mass scale [12, 72, 74].

Broadly speaking, there are two competing approaches in the use of atomic transitions to search for LPI violation. The first approach is based on astronomy: by looking at absorption lines from various atoms and molecules in space and comparing them with those observed in the same atoms and molecules in laboratories on Earth, we can leverage huge spatial separations and look-back times to get a sensitive measure of small spatial or temporal gradients of α and μ. The results of such observations have been intriguing: absorption spectra from quasars several billion light-years away indicate spatial gradients in α following a dipole distribution, with the Keck telescope in Hawaii seeing a slightly smaller α in the distant reaches of the universe [277] but the Very Large Telescope in Chile seeing a slightly larger α [276]. Separately, spectra from ammonia within the Milky Way seem to have revealed an offset $\Delta\mu/\mu = 2.6(4) \times 10^{-8}$ between μ on Earth and μ in the interstellar medium, supporting the theory of a so-called chamelion mechanism whereby fundamental constants are predicted to vary with local matter density [155]. Meanwhile, other astronomical measurements have thrown up null results for temporal gradients in μ [13, 127] and for 'chamelion-like' variations in α [267]. However, although the astronomical results are suggestive, the considerable expense of each telescope prohibits a large number of truly independent measurements from being made – the α dipole is inferred from just two telescopes – while the complex data analysis leaves room for potential errors and the ultimate results inevitably rely on unverifiable cosmological assumptions.

A second, complementary approach in the search for possible LPI violations – and the one of most relevance to us – is based on comparisons between ultra-precise atomic clocks. The main focus has been to try to detect two possible effects: (1) linear time-variations of α, μ and m_q/Λ_{QCD} and (2) annual oscillations in those constants caused by coupling to the sun's gravitational field. More recently other possible LPI violations have also been investigated, motivated by some theories of dark matter which could cause fast oscillations [273] or transient changes [68] in α. So far, however, experiments comparing atomic clocks have all yielded null results for every type of variation. The many frequency ratios that have stayed disappointingly steady over time include Hg^+/Al^+ [224], Hg^+/Cs [83], Sr/Cs [31], Dy/Cs [146], Yb^+/Cs [91, 116], and Cs/Rb/H [96, 210]. Notable among these is the Yb^+/Cs ratio, which in combination with historical data gives the tightest constraints on time variations $\dot{\mu}/\mu = 0.1(1.2) \times 10^{-16}/\text{yr}$ and $\dot{\alpha}/\alpha = 0.7(2.1) \times 10^{-17}/\text{yr}$ [91, 116].

Since the progress in microwave atomic clock accuracy has slowed in recent years (see figure 1.3) and the timebase of published microwave frequency ratios has already extended beyond a decade [96], the tightest constraints on LPI are more likely to come from direct, high-precision comparisons between optical clocks in the future[1]. The most propitious candidate currently operational is the Yb^+ octupole optical clock, since it possesses an unusually high sensitivity to α [74] while still being amenable to reaching 10^{-18} total frequency uncertainty [114]. A series of comparisons of Yb^+ against an accurate 'reference' optical clock (such as an Sr lattice clock like the one discussed in this thesis) could soon be used to provide the most stringent probe yet for terrestrial variations in α. Alternatively, recent proposals have emerged to use clock transitions in highly-charged ions which can be more than 10 times as sensitive to α than Yb^+ octupole [22, 75, 227–229]. The realization of a highly-charged ion clock requires heroic experimental effort, but a group in Germany have recently reached an important milestone by trapping a crystal containing one Ar^{13+} ion and one Be^+ ion [237]. If the systematic shifts in the highly-charged ion can be controlled sufficiently

---

[1] Note that optical clock transitions are only sensitive to α, meaning that further measurements of the other constants μ and m_q/Λ_{QCD} will have to involve at least one microwave clock [96] or, more speculatively, the study of molecular transition frequencies [20, 53, 302].

well, then we might eventually expect this research to yield another leap forward in precision LPI measurement.

Local Lorentz Invariance: We can also use atomic clocks to investigate another cornerstone of general relativity: Local Lorentz Invariance (LLI). This states that the outcome of any freely-falling local experiment should be independent of the velocity and orientation of the apparatus, i.e. non-gravitational experiments should always follow special relativity. Experimental studies into the effects of special relativity have a rich history, starting with the Michelson-Morley interferometer in 1887 which famously failed to find evidence for the luminiferous ether [182], and later reinforced by the Kennedy-Thordike [134] and Ives-Stilwell [122] experiments in the 1930s which probed different aspects of Lorentz contraction and time dilation. These three experiments all confirmed the predictions of special relativity, and in the context of test theories in which the speed of light c is only assumed to be isotropic in one 'preferred' reference frame, they collectively impose tight constraints on all of the possible degrees of freedom of Lorentz violation[1] [222]. However, several string theories imply a spontaneous breakdown of LLI at some level [139], while one explanation of the abundance of matter over antimatter is based on a violation of CPT symmetry [220] which implies either the existence of nasty non-local interactions or (again) the breakdown of Lorentz symmetry [93]. These theories motivate the present-day search for possible tiny violations of special relativity – and just as with the LPI discussed above, if an LLI violation were detected it could help refine the search for unifying theories beyond the standard model.

One way of using atomic clocks to test LLI is to set up an atomic analogue to the Michelson-Morley interferometer. The basic idea is to prepare the atom's electron cloud into two states oriented along orthogonal directions – a state of maximum $|M_F|$ which spreads along the quantization axis $\hat{\mathbf{z}}$ and a state of minimum $|M_F|$ which spreads out in the plane orthogonal to $\hat{\mathbf{z}}$. Then, over an extended period the energies of those two orthogonal states can be compared, while constantly keeping the quantization axis fixed in the laboratory frame so that it rotates with the Earth over the course of a day. In the presence of a LLI-violating 'preferred' reference frame the energy difference between the states should change with orientation, by an analogous mechanism to the hoped-for change with orientation between the two orthogonal

---

[1] In the 19th century this preferred frame was tied up with the notion of a luminiferous ether, but plausible modern guess would be the frame in which the cosmic background radiation is isotropic in energy [236], a frame which happens to travel at 380 km s^{-1} relative to the sun.

arms of the Michelson-Morley interferometer. Because the electron moves at relativistic speeds within the atom, and because the energy difference (a frequency) can be measured very precisely, these atomic clock experiments can be very sensitive probes into Lorentz violations. Unfortunately they require electronic angular momentum in at least one of the clock states (which rules out the use of lattice clocks), but null results have been obtained at the 18th decimal place for LLI violation in Ca^+ [216] and Dy [111], and a proposal to use the Yb^+ octupole clock transition looks likely to yield significantly improved constraints in the near future [76].

We might also expect atomic clocks to be well suited to the direct measurement of time dilation (perhaps using an Ives-Stilwell-type experiment). However, so far the most precise tests of time dilation have chosen to leverage extreme speed rather than frequency precision, probing a fairly broad transition in Li^+ ions travelling at up to $0.34c$ in a storage ring [41, 220]. It hasn't yet been possible to improve on these experiments using atomic clocks, perhaps because it's quite difficult to accelerate a full-blown clock to relativistic speeds in the lab. However, this might change if the latest generation of optical atomic clocks could be launched into space!

Space missions: Several proposals have emerged in recent years to launch accurate atomic clocks into space, potentially enabling a new set of science experiments beyond the capabilities of ground-based clocks. One ambitious example is the Search for Anomalous Gravitation using Atomic Sensors (SAGAS) in 2007 [289], which aimed to launch a probe into the outer solar system with a payload containing three modules: an optical atomic clock, an atom interferometer, and a laser link for communication with ground stations. The on-board atom interferometer was intended to detect any non-gravitational accelerations that would otherwise skew the data from the optical clock, while the clock could then map out the effects of both special and general relativity using continuous frequency comparisons with an optical clock on Earth. The proposal was not approved, mostly because the technology wasn't considered sufficiently developed at the time, but its science goals make for interesting reading:

- To measure time dilation at high probe velocities with a view to detecting possible violation of Local Lorentz Invariance
- To search for possible gravitational coupling of fundamental constants in violation of Local Position Invariance
- To map out the gravitational redshift from the Kuiper belt, hoping to shed light on the long-standing 'mass problem' in the formation of large Kuiper Belt Objects

- To measure the Shapiro time delay during occultation by the sun as a constraint on relativity-violating parameters within Parametrized Post-Newtonian gravity
- To verify the anomalous acceleration results from the Pioneer 10/11 probe [10], looking for further evidence of anomalous long-range gravity

A contemporary of SAGAS was the Einstein Gravity Explorer (EGE) proposal [236], which hoped to launch two atomic clocks – one optical and one microwave – into Earth orbit on a shared satellite. For reasons of technological unreadiness the EGE, too, was rejected. Many of the science goals were similar to those listed above, but a notable addition was to facilitate improved inter-continental comparisons between terrestrial atomic clocks via the EGE satellite, enabling further tests of relativity as well as improved evaluations of the earth gravitational potential.

Further to the rejected proposals, there have also been a handful of successful bids which are now bringing the concept of accurate space-clocks much closer to reality. An example is the Space Optical Clocks (SOC) program, which hopes eventually to put a lattice clock onto the International Space Station (ISS) [235]. As an intermediate step, the SOC consortium has recently realized a portable optical lattice clock prototype (SOC-2) [39] to which our group at NPL contributed a compact permanent-magnet Zeeman slower [107]. Meanwhile the Atomic Clock Ensemble in Space (ACES) project has developed an impressively robust pair of microwave clocks, one H-maser and one Cs beam clock, soon expected to bring state-of-the-art microwave clock accuracy and stability to the ISS [49]. The ACES clocks aren't quite as high-performance as the best optical clocks, but they are still expected to yield competitive tests of LPI violation as well as facilitating an order of magnitude improvement in satellite-based intercontinental frequency comparisons.

These active space programs are an important step towards reaching the lofty science goals of SAGAS and EGE, but they also contribute to a broader theme of research. Many of the technologies developed for space clocks will be needed for missions like the Matter Wave eXplorer of Gravity (MWXG) [77] and the SpaceTime Explorer and Quantum Equivalence Space Test (STE-QUEST) [1, 9], which both aim to send atom interferometers into space to test the Weak Equivalence Principle (WEP) of general relativity. More generally, we can learn from the development of space-qualified lasers, vacuum chambers and electronics for missions like SOC and ACES to expedite the design of robust equipment for use in terrestrial atomic physics laboratories and field-deployable quantum sensors.

1.2.2.2 Technological applications

The truth is that nobody knows exactly what technologies might develop from the latest generation of atomic clocks. By analogy, we can look back to the 1950s when the first Cs atomic clocks were being built by Essen and Parry [78] – at that time no-one could have anticipated the emergence, decades later, of an atomic-clock-based Global Navigation Satellite System (GNSS) as ubiquitous as GPS.

Nonetheless, we can make some educated guesses about the future. For example, GNSS systems like GPS could benefit from the improved stability of optical atomic clocks, leading to faster acquisition of accurate position information on the ground and much longer periods between re-synchronization of the clocks up in orbit. Moreover, more accurate optical clocks in orbit could be used in communication with ground stations for the many research applications that currently use GPS as a frequency reference, and could also be exploited to establish a global time standard largely free from the nefarious effects of gravity variations at ground level [236, 290]. Thinking further into the future, a GNSS network of orbiting optical clocks could enable accurate positioning not just on Earth but into deep space as well, conceivably guiding asteroid mining expeditions [47] or deep-space science probes with a precision well below 1 m. Along different lines, we can expect optical clocks to enhance the technique in astronomy of Very Long Baseline Interferometry (VLBI), which requires precise timing information to extract high-resolution radio images of distant objects so that we can learn more about the formation of stars [200].

A particularly promising application of atomic clocks is relativistic geodesy. Whilst accurate atom-interferometer gravimeters can take precise measurements of the gravitational acceleration $\vec{g}$ [62], optical atomic clocks are instead sensitive to a relativistic redshift proportional to the gravity potential U, where $\vec{g} = \vec{\nabla} U$. This effect can be used to measure elevation with surprising precision – on Earth's surface, the gravitational redshift is around 1×10^{-16}/m, and so a clock with 10^{-18} accuracy should be able to determine height above the geoid to a 1 cm level (see e.g. the experiment in [56]). Using atomic clocks and gravimeters together, it should therefore be possible to detect a wide range of subterranean events which have subtle effects on ground elevation and density, such as the depletion of aquifers, tidal or seismic deformation of the Earth's crust, and the inflation of magma chambers prior to volcanic eruption [37, 38]. If these precursors to natural disasters could be successfully detected with atomic sensors, it could have a dramatic impact on human welfare in several geologically active regions around the world.

1.2.2.3 Atomic physics

It is quite remarkable how much progress has been made in the manipulation of atomic systems in the past few decades, whether for making useful atomic sensors or for the observation of exotic quantum effects like non-local entanglement [225] or Bose-Einstein condensation [11]. The prospect of further progress in experimental atomic physics is one important driver for continued research into improved atomic clocks – what new and interesting behaviours can we create in the atomic reference sample, and how can they be harnessed to make better clocks?

For example, in one branch of proposals the aim is to use entangled, spin-squeezed states to reduce the measurement noise in atomic sensors below what is possible using classical, statistically independent atoms [138, 287]. For atomic clocks in particular, this lower measurement noise could dramatically improve the ultimate clock stability by a factor up to $\sqrt{N}$, where N is the number of atoms used per clock cycle. Towards this goal, several methods have been proposed to apply a spin-squeezing Hamiltonian in lattice clocks – e.g. Rydberg dressing [88], non-demolition measurement [142, 181], or transport of an auxiliary entangling atom [281] – though none have yet been successfully implemented. These methods could spur on another great leap forward in clock performance in a few years time, though as we discuss in section 3.4 there may be other technical obstacles to overcome first.

Another group of interesting proposals intend to create exceptionally stable lasers using cavity quantum electrodynamics (cQED) systems, where the clock reference atoms are enclosed in a high-finesse optical cavity. One such example is the superradiant laser, where a lattice-trapped cloud of atoms is continuously pumped into the upper clock state to create a population inversion (just as in a normal laser) so that the atoms radiate on the clock transition into a nearby cavity mode. Through cavity-mediated interactions, the atomic dipoles automatically synchronise into a Dicke superradiant state, thereby collectively emitting laser light at a higher-than-normal rate $\propto N^2$ rather than $\propto N$ [70]. Nonetheless the clock transition is very weak, and a realistic number of atoms is only on the scale of a few million, so the superradiant laser is expected to emit less than a pW of power. However, this small amount of laser light ought to exhibit exceptional spectral purity even below the mHz linewidth of the clock transition itself (see the original proposal [181]), enabling phase-locking of another laser with unprecedented frequency stability. Initial steps towards the superradiant laser are documented in [35, 36] (Rb Raman) and [198, 199] (Sr). An alternative cQED proposal is outlined in [177] and experimentally investigated in [285] – in this case, the cavity-enclosed atoms are used as a more passive reference, exploiting an effect of bistability to create a highly nonlinear frequency response. Like the superradiant laser, the fundamental limits of such cavity-enhanced spectroscopy are modelled to be at the sub-mHz level.

To take a wider view, optical lattice clock systems could potentially be used as a platform for atomic physics experiments beyond the field of clock-making. For example, by probing atoms trapped near the retro-reflecting mirror of an optical lattice trap, we might resolve the different regimes of atom-wall interaction – van-der-Waals, Casimir, and Lifshitz – which dominate in different lattice sites at increasing distance from the mirror [67]. Further, lattice clocks with fermionic atoms (like ^{87}Sr) can also be adapted to observe SU(N)-symmetric atom-atom interactions, allowing us to probe models of quantum magnetism [176, 221, 304].

1.3 Overview of this thesis

We begin in chapter 2 by introducing the neutral strontium atom, and describing the apparatus used in our lab to capture, cool, and trap it. In chapter 3 we then describe our clock's local oscillator – the clock laser – and the procedure used to lock it to the strontium clock transition. We discuss in detail some of the main causes of instability in the lattice clock, including detection, quantum projection, and Dick-effect noise. Finally in chapter 4 we focus on the accuracy of our lattice clock, presenting a systematic evaluation of the ^{87}Sr clock and exploring prospects for a novel high-accuracy ^{88}Sr lattice clock.

Chapter 2

Trapping cold strontium

In this chapter we describe the largest part of our strontium optical lattice clock: the apparatus needed to capture, cool, and trap samples of neutral strontium atoms.

2.1 The strontium atom

Strontium (Sr) was the first element proposed as the basis for a lattice clock [128], and since then it has become the most widespread species of optical atomic clock in the world [3, 34, 39, 81, 144, 159, 165, 212, 250, 259, 271, 293][1]. With several institutions helping to develop and characterise the technology, as of 2015 the strontium optical lattice clock has become the most accurate and the most stable frequency standard ever built [193].

What makes Sr such a good candidate for atomic clocks? The main advantage is its electronic structure, depicted in figure 2.1. Two near-cycling transitions at 461 nm and 689 nm facilitate efficient laser cooling to temperatures below 1 µK (see section 2.3). Meanwhile, a third transition at 698 nm is almost ideal as a clock transition – it is highly forbidden, enabling the resolution of very narrow lines (see chapter 3), and both clock states have zero electronic angular momentum, reducing vulnerability to Zeeman and non-scalar Stark shifts (see chapter 4). It also benefits from convenient repumping wavelengths at 679 nm and 707 nm, required for better cooling and for normalised detection of clock excitation fraction (see sections 2.3.1.1 and 2.2.4.2). Finally, though not shown on the term diagram, Sr possesses a handy 'magic wavelength' at 813 nm where the net Stark shift to the clock transition is zero

---

[1] The twelve cited institutions with Sr lattice clocks are: the University of Tokyo [259], JILA, Colarado [165], SYRTE, Paris [144], PTB, Germany [81], NICT, Tokyo [293], LENS, Florence [212], NMIJ, Tokyo [3], RIKEN, Tokyo [271], Nicolaus Copernicus University, Poland [34], NIM, China [159], the University of Birmingham, UK [39], and VNIIFTRI, Russia [250].

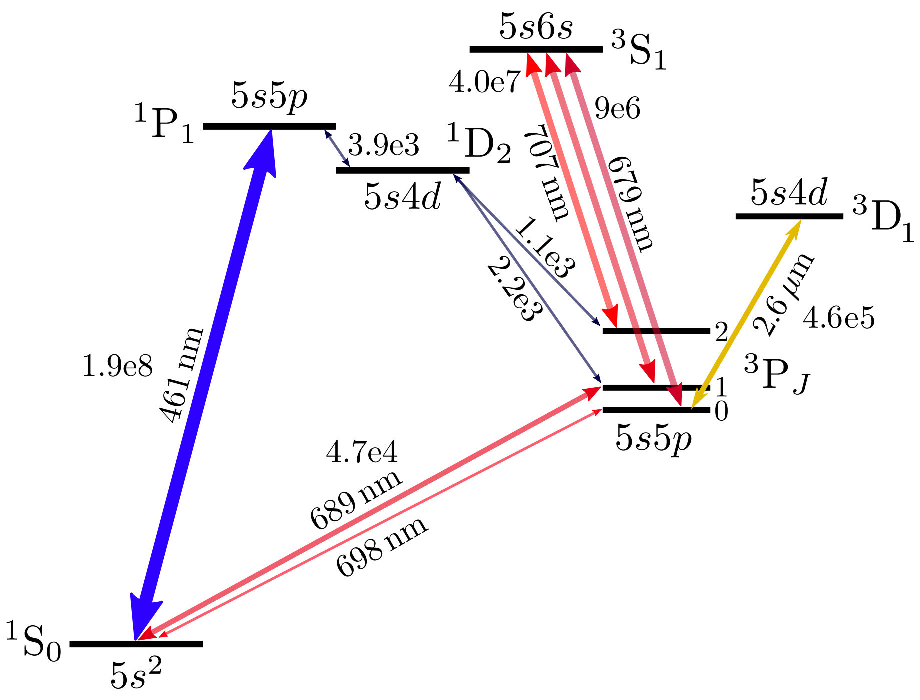

Figure 2.1: Simplified term diagram for neutral strontium. Decay rates are shown in units of s^{-1}. Hyperfine structure (^{87}Sr only) is not shown.

(see section 2.3.3.3), allowing us to confine the Sr atoms in an 813 nm lattice trap for long periods of interrogation by the 698 nm clock local oscillator.

These features are what make strontium optical lattice clocks possible in theory, but an important practical factor has been the availability of suitable laser technology – all the required wavelengths can be conveniently accessed using semiconductor sources, and many are also in range of tunable Ti:Sapphire lasers. Perhaps this is why, at least for now, progress with Sr has outpaced the other viable types of lattice clock listed in table 2.1.

2.1.1 Isotopes of strontium

From the perspective of clock-makers, two isotopes of strontium are of the most interest: ^{88}Sr and ^{87}Sr. Our clock system can work with either of these isotopes, with the ability to switch between them within just a few minutes of twiddling with equipment in the lab.

^{88}Sr is the most common isotope with 82.58% natural abundance, and is also the easiest to capture and cool – it has zero nuclear spin, and therefore none of the cooling complications that arise from hyperfine structure. However, because of the

Table 2.1: Neutral atomic species for optical lattice clocks

Species[a]	Clock	1st-stage cooling		2nd-stage cooling		Lattice	Status
	λ (nm)	λ (nm)	γ (MHz)	λ (nm)	γ (kHz)	λ (nm)	
Be [187]	455	235	88	-	-	300	Theory
Mg [85, 141]	458	285	78	383	25000[b]	468	Experiment
Ca [27, 67]	660	423	34	657	0.4	739	Mix[c]
Sr [259]	698	461	30	689	7	813	Experiment
Zn [67, 275]	309	214	112	308	4.1	416	Theory
Cd [45, 67]	332	229	91	326	65	419	Mix
Hg [178, 294]	266	254	1.3	-	-	363	Experiment
Yb [15]	579	399	28	556	182	759	Experiment

---

[a] We list group-II atoms and Yb since their $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_0$ clock transitions suppress vector/tensor lattice shifts. We exclude Ba and Ra because their lowest-lying $^3\mathrm{D}_1$ manifolds quench $^3\mathrm{P}_0$ [64, 97].

[b] This 2nd stage is operated on the $^3\mathrm{P}_2 \rightarrow {}^3\mathrm{D}_3$ line, enabling sub-Doppler cooling [179].

[c] Quench-cooled Ca [27, 63] has been used in ballistic clocks [249] but not yet in a lattice clock.

$J = 0 \nrightarrow J = 0$ selection rule, the $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_0$ clock transition in ^{88}Sr is completely forbidden to all orders in the multipole expansion. The positive consequence of this is that the natural linewidth of the transition is exceptionally narrow – the dominant decay route becomes a two-photon E1M1 process taking an estimated 5800 yrs [232]. The negative consequence is simply that it is impossible to excite the clock transition with a single photon, which plainly undermines efforts to use the transition as a frequency reference. The solution is to apply at least two photons, usually including one 'DC photon' comprising a static magnetic field [254], but this approach requires strong fields which induce large systematic Stark and Zeeman frequency shifts (see section 4.2). A further disadvantage of ^{88}Sr (a boson) is the presence of s-wave collisions, which can induce significant decoherence and frequency shifts of the clock transition [161]. As a result of these problems, most research groups have turned to the odd ^{87}Sr isotope instead.

^{87}Sr is quite a bit less common with only 7.00% natural abundance, but its non-zero nuclear spin $I = 9/2$ lets us avoid the biggest shifts plaguing ^{88}Sr. The hyperfine interaction weakly mixes $^3\mathrm{P}_0$ with the $^3\mathrm{P}_1$ and $^1\mathrm{P}_1$ manifolds, allowing single-photon $^3\mathrm{P}_0 \rightarrow {}^1\mathrm{S}_0$ decay at a rate of around $\Gamma/2\pi = 1\,\mathrm{mHz}$ [43] – this lets us excite the transition in reasonable time even with very weak probes, but still means that we can resolve a respectably narrow linewidth. Further, s-wave collisions between identical fermionic ^{87}Sr atoms are forbidden by symmetry, dramatically suppressing collisional decoherence and frequency shifts. These features have led to the dominance of ^{87}Sr as the isotope of choice when building super-accurate lattice clocks [81, 145, 193, 271].

However, the non-zero nuclear spin in ^{87}Sr is no panacea. The resulting messy Zeeman structure complicates the cooling process, necessitates an extra state-preparation stage, and exposes the clock transition to substantial 4.9 kHz/mT linear Zeeman shifts as well as tensor and vector Stark shifts from the lattice trapping light (see section 4.1). Partly for these reasons we revisit the forlorn ^{88}Sr isotope in section 4.2, exploring its potential as a magnetically insensitive, highly accurate clock.

2.2 Hubs of the Sr experiment

In this section we introduce some of the central elements of our strontium lab. These systems have turned out to be reliable workhorses when running the experiment for extended periods, underlying more than anything else our ability to run consistently over several-weeks-long measurement campaigns such as ITOC in 2015 – even though it took rather a lot of work, it was the right choice to build it right rather than in a hurry! Here, we focus on:

- The transfer cavity, which we use to peg the frequencies of all our cooling and trapping lasers to the ultrastable clock laser
- The FPGA-based experimental control system coordinating the clock cycle
- The ubiquitous servo electronics for intensity, temperature and current control
- The ultra-high vacuum science chamber in which the cold strontium is prepared

2.2.1 The transfer cavity

The purpose of the transfer cavity is to provide a series of very stable resonances to which we can lock the frequencies of the many lasers we use for cooling and trapping strontium (689 nm, 922 nm, 813 nm, and sometimes 679 nm – we discuss these lasers and their application in 2.3). At the cavity resonances, constructive interference makes the ordinarily highly-reflective mirrors suddenly start transmitting a large fraction of the incident light (see the transmission data in figure 2.3); by steering our cooling lasers towards the peaks of the cavity resonances, we can keep them from drifting away from their respective atomic transitions.

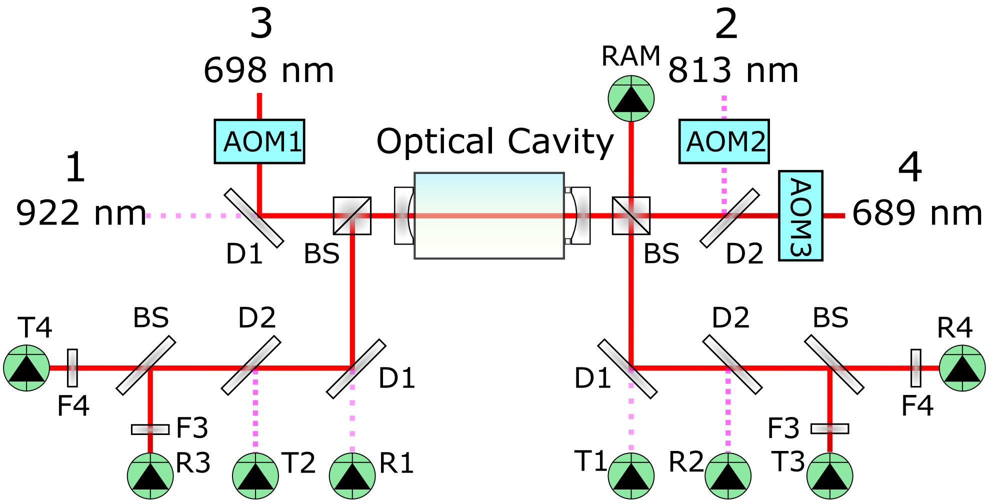

Figure 2.2: The optical setup surrounding the transfer cavity. This apparatus was originally conceived by Ian Hill, as described in his thesis [106]; here, I present the details of its most recent incarnation. D1 and D2 are Thorlabs DMLP900 and DMLP605 dichroic beamsplitters, used to separate the 922 nm and 813 nm beams from all the shorter wavelengths. Meanwhile, F3 and F4 are optical bandpass filters (Thorlabs FB700-10 and FB690-10), while the various AOMs are required to bridge frequency gaps between the cavity resonances and the desired laser frequencies.

The frequencies where the cavity resonates can be written as [231]

$$\nu_{nlm} = n\nu_F + (l+m+1)\frac{\Delta\xi}{\pi}\nu_F \tag{2.1}$$

where $\nu_F = c/2L = 1.93\,\text{GHz}$ is the cavity free spectral range, n indexes the longitudinal modes (roughly how many half-wavelengths fit inside the cavity), (l, m) indexes the transverse Hermite-Gaussian modes (typically $(0, 0)$), and $\Delta\xi = \xi(x_1) - \xi(x_2)$ is the difference in Guoy phase at the positions of the two cavity mirrors (unimportant except for lifting degeneracy between transverse modes).

To a good approximation the resonant frequencies $\nu_{nlm} \approx nc/2L$ are determined by the cavity length L, underlying the mechanism by which the transfer cavity works: (1) We stabilise the length of the cavity to the frequency of the ultrastable clock laser using the piezo mounted to one of the cavity mirrors, then (2) we use the (now fixed) cavity modes to stabilise the frequencies of all the other lasers in the lab.

However, there are some nefarious effects which undermine the transferred frequency stability. For instance, dispersion in the refractive index $n(\nu)$ inside the cavity causes the optical path length $L(\nu)$ between the mirrors to be frequency-dependent. If this dispersion changes over time (from e.g. changes in air pressure), then the

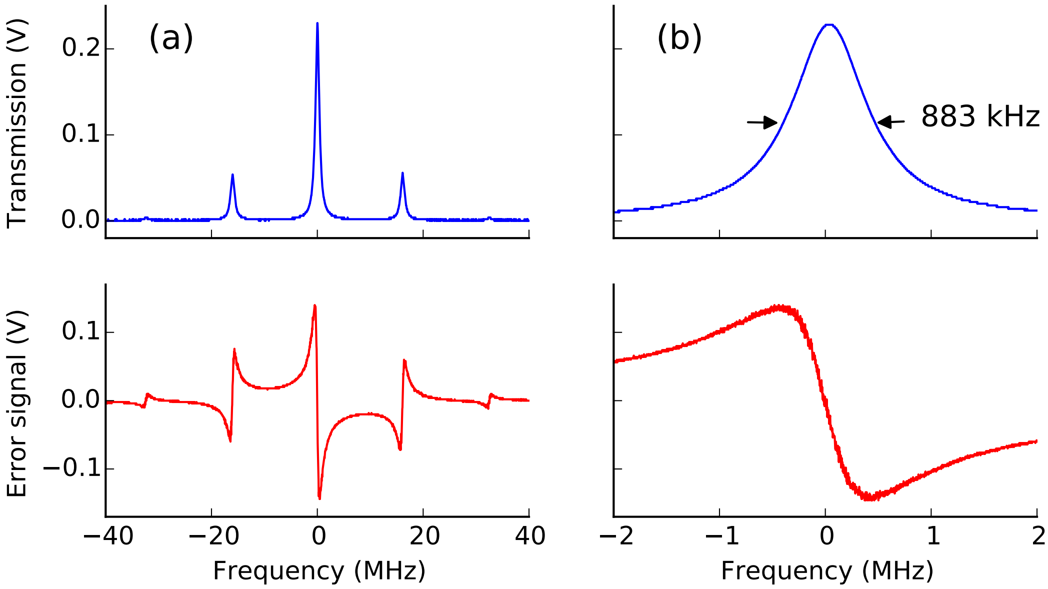

Figure 2.3: The transmission and PDH error signal from the 689 nm laser scanned over resonance of the transfer cavity. (a) shows first- and second-order sidebands from the 16.05 MHz phase modulation, while (b) zooms in on the central feature. 16 scans are averaged with the 689 nm laser locked to a fixed pre-stabilisation cavity.

resonant frequencies ν_{nlm} change independently of each other in a way that can't be cancelled out using the piezo. We prevent variations in $n(\nu)$ in our setup by placing the cavity under vacuum – and indeed, when the ion pump failed, slowly letting air leak into the cavity, we observed degraded transfer stability until the pump was eventually replaced.

After dispersion, the factor that ultimately limits our transfer stability is imperfections in the Pound-Drever-Hall stabilisation process.

2.2.1.1 Pound-Drever-Hall locking

To frequency-stabilise all the lasers to their nearest cavity resonances we use the Pound-Drever-Hall (PDH) technique. There is plenty of literature on PDH locking, starting with the original paper [73] and including detailed hands-on guides such as the one given in [84]. In our transfer-cavity implementation, after going through a double-passed AOM to bridge to the nearest cavity resonance, each laser beam is sent through a bulk electro-optic modulator (EOM) to apply phase-modulation, and then the reflection from the cavity is picked off onto a fast photodetector (see figure 2.2). Demodulating the photodetector signal to DC (i.e. sending it into a double-balanced mixer with a split-off version of the EOM drive signal), we realise an 'error signal' shaped like the one in figure 2.3. Qualitatively, the shape of the error signal can be explained as follows:

- When far from resonance the carrier and sidebands all gather the same 180 degree phase shift on reflection, so the total beam remains purely phase-modulated with no amplitude modulation – therefore no beat is present on the photodetector and the error signal is zero.
- When the carrier (or a sideband) is near resonance, the phase of that carrier (or sideband) is changed by an amount different from 180 degrees on reflection – therefore it starts to form an amplitude beat with the other sidebands on the photodetector
- In particular, when the laser is above resonance, the resonating light circulating the cavity *lags* in phase relative to the driving light. When it's below resonance, the circulating light *leads* in phase. The net result, once the cavity leakage interferes with the directly reflected light, is that the error signal is positive below resonance and negative above resonance – exactly the kind of signal we can use to steer the laser towards the cavity mode.

Of course, all of these features of the PDH error signal can be confirmed analytically using some straightforward mathematics – see for example [30].

2.2.1.2 Residual amplitude modulation (RAM) servo

PDH locking is remarkably effective at forcing a laser's frequency to track the cavity length – the method can often combine high lock bandwidth with excellent DC accuracy. No system is perfect, however, and care must be taken to prevent potentially harmful offsets from emerging.

With careful design and implementation, electrical offsets and instabilities in the PDH lock can usually be kept at a negligible level. Precautions include mixing the photodetector beat with the correct phase, having sufficient power into the LO and RF mixer ports, using low-drift op-amps and stable voltage references in the loop filter, shielding all signals properly to avoid pickup, grounding things sensibly, and so on. After the electronics, the dominant source of lock offset – and therefore laser frequency instability – is usually optical. The mechanism for optical offsets is residual amplitude modulation (RAM): any element of amplitude modulation that leaks onto the beam incident on the cavity will look exactly like a non-zero error signal on the PDH photodetector.

RAM can arise from several sources. For instance, EOM crystal birefringence can lead to accidental modulation of the polarisation angle which gets converted to RAM further down the beam path by polarisation-selective optics. The birefringence itself depends on temperature, DC bias voltage, and stresses and defects in the crystal [291] – all of these are therefore potential sources of RAM drift (but they can potentially also be used as tools for RAM stabilisation [156, 291, 303]). To suppress EOM-sourced RAM we have found it helpful to avoid having any polarisation-selective optics after the EOM wherever possible. Further intensity-dependent RAM can occur from photorefractive losses in the EOM if they attenuate the PDH frequency components by different amounts [234]. Finally, even if the beam coming out of the EOM has perfect phase modulation, parasitic etalons further down the beam path can change the relative phases of the carrier and sidebands – a telltale sign of parasitic etalons is when slow error-signal fluctuations are observed by scanning the laser back and forth by a few GHz; the solution is typically to use high-extinction optical isolation and lay out all the optics in a carefully skewed orientation.

For reasons we never fully diagnosed, the polarisation modulation from our 689 nm EOM is particularly large and unstable. The effect of the polarisation modulation is compounded by the need to fibre-deliver the light to the transfer cavity – the laser is too far from the cavity for free-space delivery: in order to avoid parasitic etalons with the fibre end face we had to install a polarisation-selective optical isolator after the fibre output, inevitably converting the polarisation rotations into amplitude modulation. RAM is particularly dangerous to us at 689 nm because our transfer cavity feature is rather wide at 883 kHz and we require a frequency stability of ~ 3 kHz in order to maintain fairly constant ^{87}Sr atom number in the lattice (see section 2.3.2).

The best way to avoid this problem would be to ensure a higher cavity finesse, ideally around 30000 within the 689-698 nm range, so that the cavity features could be much narrower. However, given the system we already had in place, the quicker solution was to implement the AOM-based RAM servo depicted in figure 2.4. In this scheme, the light incident on the cavity is picked off onto a separate photodetector, and both the in-quadrature (Q) and the in-phase (I) components of RAM are mixed down to DC. Using voltage-controlled attenuator setups that can pass positive- or negative-amplitude RF signals, we can flexibly apply I and Q amplitude modulation to the carrier of an AOM in the 689 nm beam-path, diverting a modulated amount of light away from the zero-order beam. We note that the optical fibre after the AOM is an essential component – it spatially filters the beam, homogenizing the AM applied by the AOM. Closing the loop, the AM applied to the AOM will exactly compensate

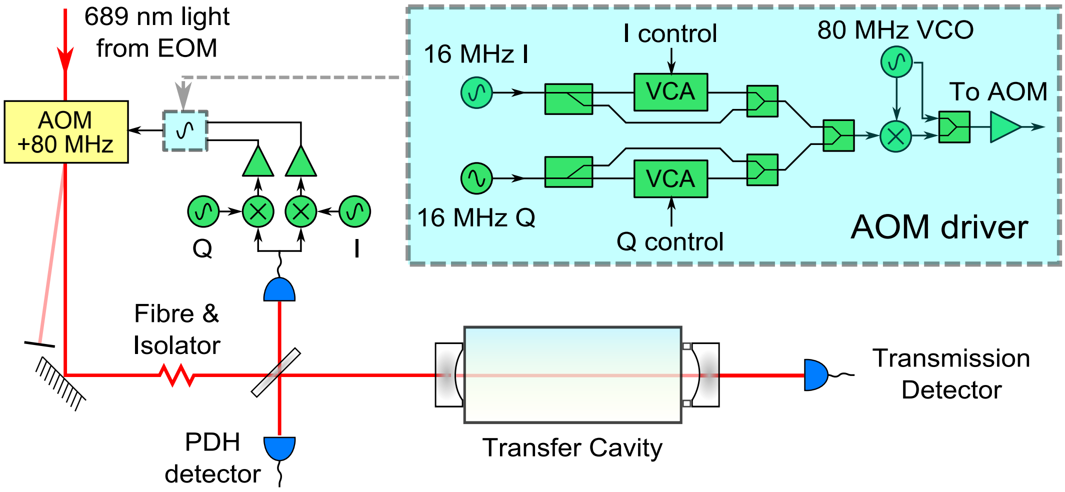

Figure 2.4: Key optics and electronics behind the AOM-based RAM servo.

the RAM from the EOM, preventing offsets in the PDH lock. We observe good RAM suppression from typically between −50 and −60 dBm to below −80 dBm on our photodetector with a lock bandwidth of a few kHz. As it turns out, John Hall (Nobel prize-winner, and the 'H' in PDH) recently built a remarkably similar 'RAM-buster' system at JILA – it was a little surreal to watch him present his experiment in a conference talk while the many similarities slowly unfolded between his setup and ours [102].

With the RAM servo in place, the transfer cavity reliably gives us excellent performance: the 689 nm system is tamed to less than 2 kHz long-term deviations, the 813 nm is kept easily within 100 kHz, and the 922 nm is controlled at the sub-MHz level. If we ever had to build another self-contained Sr system, we would therefore choose the same type of transfer cavity scheme, in stark preference to using e.g. independent saturation-spectroscopy-stabilised lasers.

2.2.2 Experimental control

Optical lattice clocks, like many atomic physics experiments, operate by continuously repeating a complex, precisely-timed sequence of events (see section 2.3). To coordinate this sequence, we need a centralised 'experimental control' system – i.e. a flexible, programmable waveform generator with dozens of output channels – capable of dynamically controlling all the different experimental parameters.

Our biggest component of experimental control, which we'll call the 'main sequence control' unit, supplies the lattice clock with 16 dynamic analogue signals for

Table 2.2: Experimental control components

Component	Hardware	Software
Main sequence control	Home-built FPGA boards	Labscript
Camera DAQ and clock servo to Sr	iXon Ultra, DDS & NI USB 6009[a]	LabVIEW
Wavemeter locks for repump lasers	WSU-10 & NI 9264	LabVIEW
Oscilloscope for remote monitoring	8-channel PicoScope 4824	PicoScope

---

[a] The NI USB 6009 data acquisition (DAQ) module is needed for running multiple servos; the 'main sequence control' program determines which servo cycle is currently underway, and communicates this information to the 'clock servo' program using 4 bits sent into the USB 6009 inputs.

ramping laser intensities, laser frequencies and coil currents, plus around 35 digital signals for e.g. shuttering laser beams and triggering the camera. In this section we outline how this home-built hardware works, and we describe how its output sequence is programmed using the Python-based Labscript suite [241].

In addition to the main sequence control unit, the lattice clock also relies on several independent auxiliary components of experimental control listed in table 2.2. The most complex of these is the 'Camera DAQ and clock servo to Sr' unit, which gathers data from the iXon camera, processes it into an error signal, and then feeds back to the clock laser frequency to steer it towards the Sr clock transition (see section 2.2.4.2 for details of the imaging process, and section 3.2.1 for a high-level description of the clock servo). Since I did not write the LabVIEW code behind this clock servo, I leave a more detailed explanation to one of its authors, William Bowden, to be written up in his doctoral thesis.

2.2.2.1 FPGA hardware

In the early days of the lattice clock experiment we chose a fairly standard solution to the problem of experimental control, buying all of our hardware (PXI M series) and software (LabVIEW) from National Instruments. However, as we kept integrating more and more lasers and cooling stages into the experiment, the control system necessarily had to grow in complexity. By late 2014, the bloated graphical LabVIEW software had made it very laborious to program the experimental sequence, and it was clear that expensive upgrades to the NI hardware were going to be needed in order to complete the lattice clock.

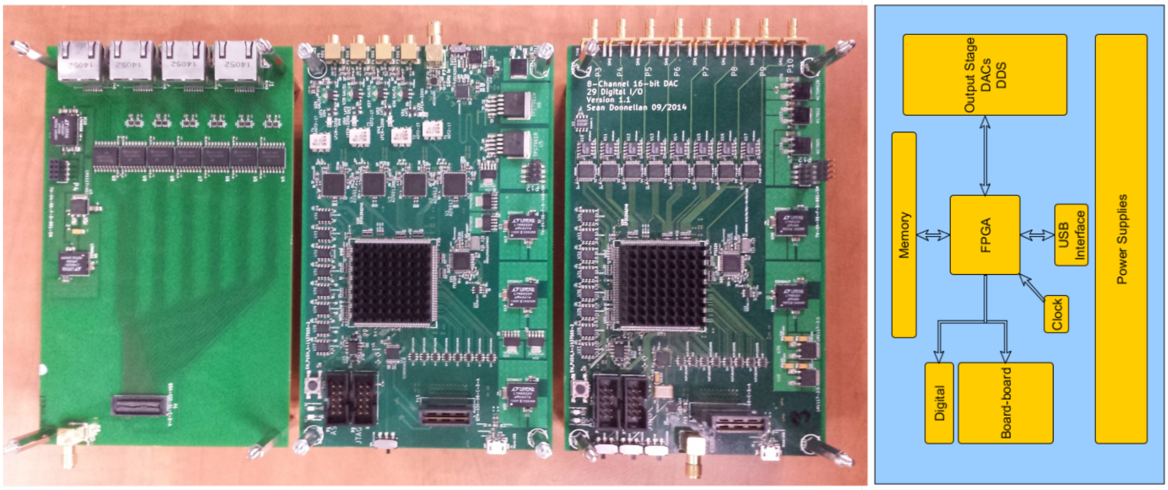

Figure 2.5: *From left to right:* (1) The header board for distribution and isolation of 24 digital outputs from the DAC FPGA board. (2) The 4-channel DDS FPGA board. (3) The 8-channel DAC FPGA board. (4) A schematic of both FPGA boards. Courtesy of Sean Donnellan.

Instead of spending lots of money to extend our already-messy NI system, we chose to replace it entirely with the new experimental control hardware shown in figures 2.5 and 2.6. These PCBs were designed in-house for our project by Sean Donnellan, so we had the freedom to: (1) match their specifications to our experimental requirements, (2) purchase several boards cheaply at production cost, and (3) tailor the hardware for integration with the Labscipt suite, an open-source, Python-based sequence-building program developed at Monash University specifically for cold atom applications. This last advantage is perhaps the most important – as we will see in the next section (2.2.2.2), the Labscript suite enables us to write complex, flexible, and extensible experimental sequences very naturally.

Each PCB in the new control system is centred around an Altera Cyclone III EP3C25Q240C8N Field-Programmable Gate Array (FPGA) (see figure 2.5). The FPGA chip contains a large number of logic elements for highly parallel signal processing, plus some phase-locked loops, multipliers, and 66 × 9 kbit chunks of embedded SRAM. As shown in figure 2.5, the FPGA is surrounded by several external elements, including:

- A bank of eight 1 Mbit SRAM modules (23LC1024) for storing sequence data
- An FTDI FT232H chip to interface between the PC and the FPGA

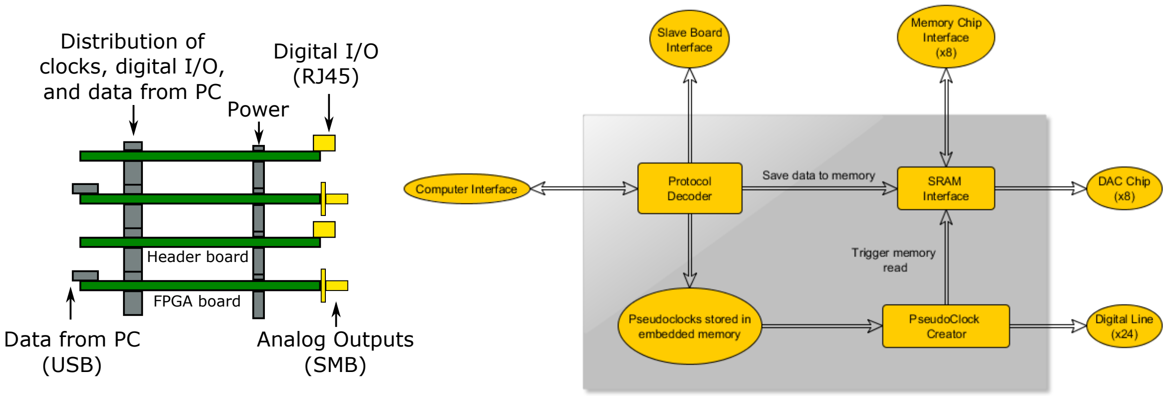

Figure 2.6: *Left:* A depiction of how the DAC-output control boards stack together with isolating header boards. This stack has 16 analogue and 48 digital output channels in total. *Right:* Schematic of the Verilog modules programmed onto the FPGA (grey box) and how they communicate with other parts of the experimental control hardware. Courtesy of Sean Donnellan.

- A bank of output chips depending on board type:
  - **DAC type:** Eight LTC1592 DAC chips, with 16-bit resolution and a settling time of 2 µs
  - **DDS type:** Four AD9912 DDS chips, clocked at 1 GHz, buffered, filtered, and amplified to produce single-ended ∼ 13 dBm outputs

The PCBs can be stacked as shown in figure 2.6 to increase the number and types of output channels available. For the work in this thesis, only two DAC FPGA boards and two header boards were incorporated into the stack – we found that this supplied sufficient channels for running a single Sr lattice clock – but there are plans to add further DAC and DDS boards to the stack in the near future. When stacking multiple boards, one board is designated the “master”; this board is responsible for communicating with the PC and for bussing information and triggers to any other boards in the stack (the “slaves”).

The working principles of the control hardware are best described by running through the process of turning it on. After the FPGA has been programmed using Sean Donnellan’s Verilog code (this program is non-volatile, so can be implemented before installation), the steps are:

1. The sequence data is compiled on the PC using BLACS, one of the programs in the Labscript suite. In order to compress the data, a pseudo-clock is included

for each channel, encoding the points in time when that output channel will change during the sequence. This allows us to specify the sequence with a timing resolution down to a single 22 ns FPGA clock cycle, even for sequences lasting several seconds, while only needing to send a few kbits per channel to the FPGA [241].

2. The compiled data is sent from the PC to the master control board via USB. The master FPGA then saves the pseudo-clock data to its internal SRAM memory, saves the DAC/DDS output values to the external SRAM chips, and saves the number of desired sequence repetitions to one of its internal logic elements. If there are any slave boards present, then all the slave-output data must be forwarded on by the master to be saved on that slave board.

3. The shot is triggered from the PC. From then on, the FPGA deterministically clocks out the saved sequence data into all the output channels in parallel, with single-clock-cycle (22 ns) resolution. The maximum modulation rates of the analogue/RF channels are, however, limited by the few-µs delay in DAC/DDS chip response.

4. Once the sequence is finished, it is repeated until either (1) the number of desired repetitions has been fully completed, or (2) a 'stop' signal is received from the PC. (Note: to allow indefinite operation of the clock, an input of 66535 desired repetitions is interpreted as infinite).

As well as acting as a complex sequence generator, the FPGA hardware also supports a 'real time' mode compatible with the Labscript 'runmanager' GUI. When we're not running the full experimental sequence, this GUI lets us manually set and toggle digital and analogue outputs at our leisure; this mode is quite handy for setting up and testing individual devices under control.

2.2.2.2 Labscript: example sequence

The various features of the open-source Labscript suite are well summarised in [241], and further documentation and downloads are available from the Labscript website. For the work in this thesis, we only make use of the Python-based labscript, BLACS, runmanager, and runviewer programs within the suite. The roles of the other Labscript programs (BIAS, lyse, and mise) are instead fulfilled by our home-written LabVIEW program, dubbed 'Camera DAQ and clock servo to Sr'.

Most of the details of the Labscript software are well beyond the scope of this thesis, but in this section we briefly outline what a front-end experimental control script actually looks like. At the top level, the script is split into three sections:

The purpose of the first section is to define objects corresponding to the FPGA hardware's output channels. Here, following a similar template as the example script on the Labscript website, we import the 'labscript' module and we write out a connection table to name all the FPGA output channels (16 analogue plus 48 digital).

The purpose of the second section is to define objects corresponding to stages in the experimental sequence. To avoid clutter in the main sequence program, a lot of the details of these objects are actually defined in a separate, auxiliary python module – this module contains a long list of class definitions, each of which is designed to generate a specific kind of stage (e.g. an `Image()` class is defined for stages where pictures are taken of the atoms). In the main program, a stage is represented by an *instance* of one of the classes defined in the auxiliary module; the parameters of that stage are then set by assigning values to the class variables. For example, to create a blue MOT stage for loading Sr (see section 2.3), we might write:

```
BlueMOTLoading = BlueMOT()
BlueMOTLoading.MOTtime = 0.55 # Seconds
BlueMOTLoading.MOTIntensity = 1.2 # Relative units
BlueMOTLoading.MOTdetuning = 415.0 # AOM frequency (MHz)
BlueMOTLoading.MOTCurrent = 75.0 # Amps
BlueMOTLoading.SBPower = 1.2 # Relative units
```

In the third and final section, the actual experimental sequence is fully laid out in an intuitive, ordered, and extensible format, combining all the pre-defined stages (class instances) listed above. We give an example below which runs four separate servos to the clock transition:

```
polDirections = [1,-1,1,-1] # Atomic polarisation (M_f) in each servo
MOTtimes = [0.8,0.8,0.08,0.08] # Blue MOT loading time in each servo
shot_reps = 65535 # Desired sequence repetitions
start()
t = 0.0
t += initialise()
for pulseNumber in [0,1]:
    for servo in range(4):
```

```
        t += sendClockServoInfo1.send(td=t, pulseNumber=pulseNumber,
                                                servoNumber=servo)
        t += BlueMOTLoading.bMOT(td=t, MOTtime = MOTtimes[servo])
        t += BlueMOTCooling.bMOT(td=t)
        t += rMOTBB.rMOT(td=t)
        t += rMOTNB_HighI.rMOT(td=t)
        t += rMOTNB_LowI.rMOT(td=t)
        t += StatePrep1.prep(td=t,polDirection=polDirections[servo]))
        t += SpecPulse6981.pulse(td=t)
        t += Image1.NormalisedDetection(td=t)
stop(t)
```

Specifically, this script was written in order to evaluate the density shift in ^{87}Sr – in section 2.3 we find a description of all the listed cooling stages, and in chapter 4 we give some context about shift evaluation techniques. However, the details of the script aren't important at this juncture[1]; the example is given here simply to highlight how the Python/Labscript software enables us straightforwardly to construct flexible, complex sequences for optimisation and evaluation of the Sr lattice clock.

2.2.3 Servo electronics

A large chunk of the hardware around the lab is centred on servo control: all the cooling lasers need to arrive at the atoms at the right frequencies and intensities, the magnetic field coils must ramp to the right current at the right times, and the clock laser beams have to be actively phase-stabilized whenever delivered via optical fibre. All in all, our Sr lattice clock system relies on more than 50 active servo loops – some in software, but most based on analogue electronics.

---

[1] For the curious I give some notes of explanation of the script here. The four servos are split into two pairs: one pair operates at large atom number by using a long blue MOT loading time of 0.8 s, while the other pair operates at low atom number by using a much shorter blue MOT loading time of 0.08 s. Within each pair, the first servo operates using atoms spin-polarised into the $M_F = +9/2$ Zeeman sublevel, while the second operates with atoms prepared into $M_F = -9/2$.

The `pulseNumber` variable indexes whether the 'high' or 'low' side of the clock transition is being interrogated (see section 3.2.1), while the `sendClockServoInfo1.send()` function is needed to send information to the 'Camera DAQ and clock servo to Sr' unit in order to tell it which servo and pulse number is currently underway – that separate program then controls the DDSs responsible for clock interrogation. Also, we observe that any class variable can optionally be overidden when the stage function is called, simply by passing that variable as a keyword argument – e.g. the length of time devoted to the blue MOT loading stage can be set to 0.8 s by writing `BlueMOTLoading.bMOT(td=t, MOTtime=0.8)`, even if we had previously set a different time such as `BlueMOTLoading.MOTtime=0.3`. This feature is especially useful for running consecutive cooling sequences with different experimental parameters.

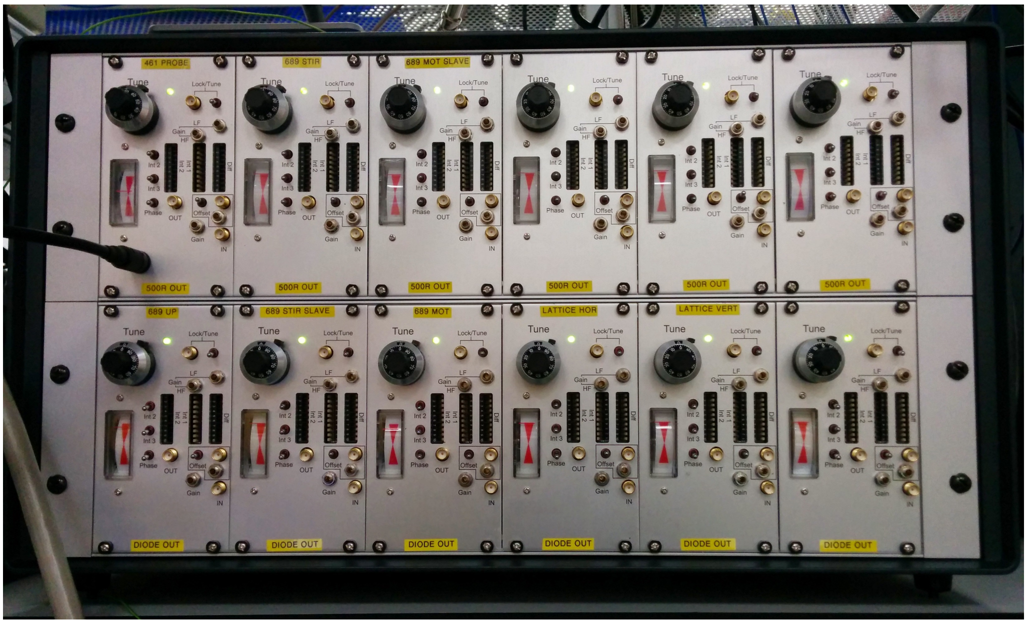

Figure 2.7: Photograph of 12 boxed-up general-purpose loop filters.

Given the volume of servo electronics required, the only affordable option was to design our own circuits for external assembly – an approach which also let us match the product exactly to our own needs. The design we use most extensively is photographed in figure 2.7. In total we had 33 loop filters of this variety made, with 24 housed in two racks of 12 plus the remainder in standalone boxes. This almost seemed like too many at the time of purchase, but we're now using all but one or two of them; the number of cables now emanating from these boxes is formidable.

Upon assembly the circuit can be configured in one of two settings: (1) as a loop filter specifically for PDH locking of ECDLs, or (2) as a general-purpose loop filter for intensity, phase, or current control. In the PDH configuration the loop-critical parts of the circuit are as shown in figure 2.8 except that 'setpoint IN' channel isn't populated and SW1 is permanently closed. The current modulation output is sent through a buffer similar to that used in [158] so that the loop filter can optionally inject current directly onto the diode in parallel with a fixed current source. The internal offset range is set at around 20 mV to make it easy to compensate small mixer offsets, while the 5 V is supplied from a REF5050 voltage reference to ensure long-term stability. All the op-amps in the signal path towards the current modulation output are chosen for their speed ($\sim$ 200 MHz GBW) as well as low noise, mostly OPA4820 except for the integrator which needs a FET input THS4631D to reduce input current noise and offset. In the 'total gain' tuning op-amp, it would have been wise to place a resistor

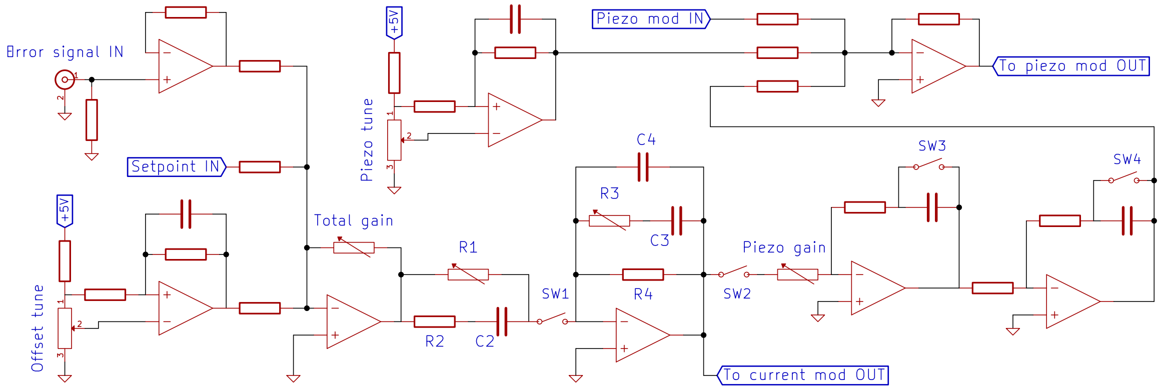

Figure 2.8: Simplified schematic of the loop filter.

between the positive input pin and ground so as to match input impedances – the fast BJT-input op-amps tend to have very large input bias currents on the scale of µA. To reduce the effect of op-amp offset currents and to achieve the highest lock bandwidth we've found it essential to use low resistor values $< 1\,\mathrm{k}\Omega$ into the op-amp input pins.

For flexible tuning of the loop filter transfer function, several controls are accessible on the front panel: the 'total gain', R1, R3, and 'piezo gain' resistors are Vishay M43-series trimmers, while C2, C3, and the integrating capacitor beneath SW3 are each selected on 10-position DIP switches. This gives us plenty of tools to optimize the gain curve for maximum bandwidth, phase margin, and DC gain (see figure 2.9). Since the cavity response has a pole at its transmission linewidth $\Delta\nu$, at the very least it is necessary to have a proportional stage (here shown at $1/R_3C_3 < \omega < 1/R_1C_3$) in order to pass unity gain with a good phase margin. To push the lock bandwidth above the MHz level we've found it essential to introduce the differentiation stage as well (i.e. to switch in a non-zero capacitance on C2). The length of the cable to the laser diode must also be as short as possible since the cable capacitance forms an unhelpful pole with the kΩ-level current-modulation output impedance. For a good in-depth discussion of PDH loop optimization see [84].

The various switches (DG417/8 series) are there to make it easy to capture and retain lock. When the lock is disengaged, SW2 is open and SW3 and SW4 are closed so that the integrators to the piezo output rest at zero – the ECDL piezo can then be freely tuned around using the 10-turn Bourns 3590 series 'piezo tune' potentiometer. In another part of the loop filter circuit (not shown), we plug the cavity transmission signal into a comparator to detect when the laser is at cavity resonance – when resonance is detected, SW2 quickly closes and SW3 and SW4 open, engaging the

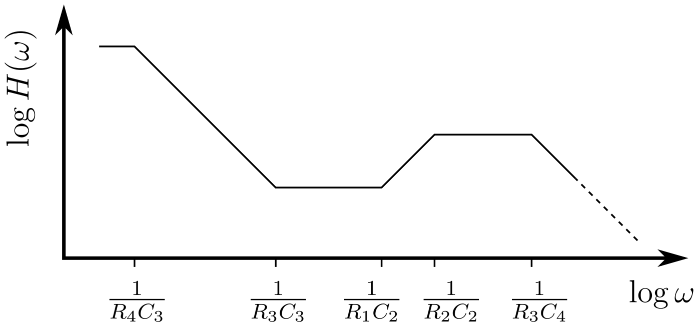

Figure 2.9: An idealized loop filter transfer function from the error signal input to the current modulation output for fast PDH locking.

double-integrated piezo lock. When the cavity transmission falls again, the switches are flipped back into ‘tune’ mode so that the laser can be scanned back onto resonance. To avoid vulnerability to transient dropouts in transmission, we introduce a delay of around 300 ms between any detected loss of lock and the integrators being reset. The lock detection can be nicely exploited by constantly feeding a slow triangle wave into ‘piezo mod IN’, so that if the laser does come out of lock it will usually recapture automatically, although sometimes a diode current tweak is also needed.

In the more numerous ‘general purpose’ loop filters we don’t populate the piezo modulation input or the current modulation output: this simpler configuration only needs to feed onto one actuator, so only the ‘piezo modulation’ output (usually not actually going to a piezo) is needed. For further simplification, SW2 is permanently closed and the redundant ‘piezo gain’ pot is replaced with a fixed resistor. Some extra features are also added to the general-purpose box: we populate a ‘setpoint IN’ channel for dynamic tuning of the lock points for e.g. intensity or current control, and we add an external digital input to switch in the loop when desired (closing SW1 and opening SW3 and SW4) – this allows us to avoid integrator windup, meaning we can capture intensity lock with minimal overshoot and within a few µs of turning on the clock probe and the fluorescence detection beams.

2.2.4 The science chamber

The ultra-high vacuum (UHV) science chamber shown in figure 2.10 is the main focus of the lab. It is inside this chamber that we generate, capture, cool, trap, and interrogate strontium atoms using a deluge of laser beams delivered from around the

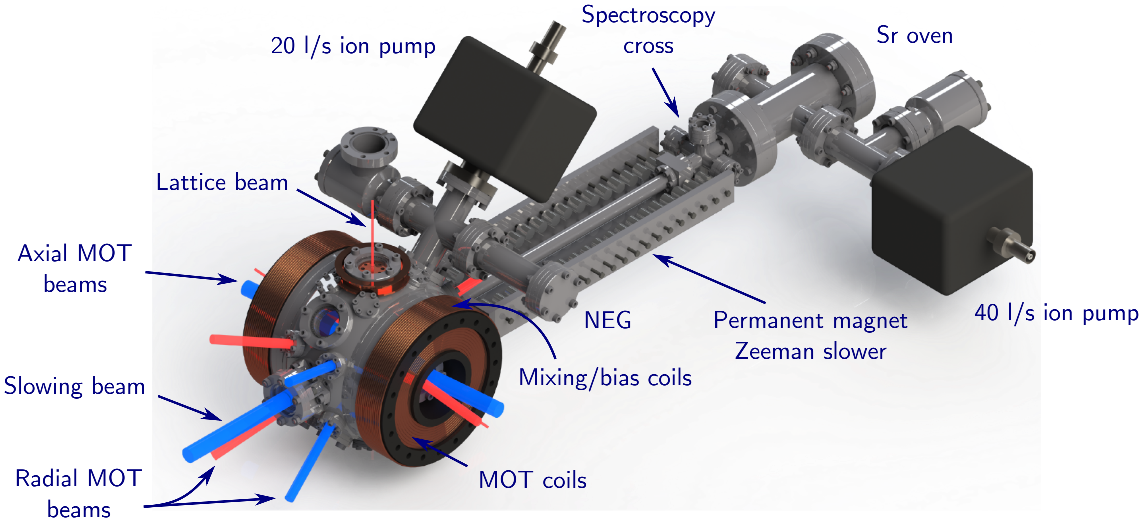

Figure 2.10: A 3D CAD model of the science chamber including some of the laser beams used for cooling and trapping Sr.

room, with the ultimate purpose of stabilizing the clock laser to the narrow $^1S_0 \rightarrow {}^3P_0$ clock transition.

The UHV chamber itself remains the same as described in previous theses [46, 106] – save for one refill of the Sr oven, it has stayed under vacuum throughout the course of my research – but its surroundings have needed to develop significantly in order to implement the long series of cooling stages required to operate an Sr clock. The main chamber is a Kimball physics MCF800-ExtOct-G2C8A16 extended spherical octagon of 8″ diameter made from 316LN non-magnetic steel, pumped by a 20 l/s Varian VacIon Plus 20 StarCell ion pump and a CapaciTorr CF 35 non-evaporable getter (NEG) from SAES Getters. We measure a consistent $1/e$ lifetime of 8 s in three independent Sr traps: (1) the quadrupole magnetic trap capturing shelved $^3P_2$ atoms leaking from the blue MOT, (2) the narrow-band red MOT (at least for saturation $s \gtrsim 100$), and (3) the magic-wavelength optical lattice trap. Assuming that these lifetimes are vacuum-limited, we have an implied background pressure of around 3×10^{-10} mbar [190].

The relatively high-pressure oven at $\sim 1 \times 10^{-7}$ mbar is separated from the main chamber by a 1 mm aperture on the axis of the atomic beam (immediately after the spectroscopy cross) and by a manually operated gate valve (after the aperture and before the Zeeman slower). The valve allows servicing of the Sr oven without breaking vacuum in the main chamber, while the aperture aids differential pumping and reduces black-body radiation leakage into the main chamber.

2.2.4.1 Magnetic field coils

As described in section 2.3, we rely heavily on magneto-optical traps (MOTs) to capture and cool the Sr atoms. For these traps to operate efficiently, we need to generate quadrupole magnetic fields with a zero in the middle of the chamber and with gradients up to 4 mT/cm which must be quickly switchable from 4 to 0.3 mT/cm within less than 2 ms. To realize this, we have bolted a pair of aptly-named 'MOT coils' onto the sides of the chamber. Each coil consists of 9x8 turns of 4 mm square-cross-section Kapton-insulated copper wire with 2.5 mm central circular bore through which ∼ 15 °C cooling water flows. The coils are potted with Stycast into a Delrin former which provides a 2 mm thick insulation layer between coil and chamber. Since they are mounted a large distance away from the atoms, the MOT coils only generate a field gradient of $0.05\,\mathrm{mT\,cm^{-1}\,A^{-1}}$, requiring a up to 80 A to be driven through them. Fortunately with water cooling the coils can operate continuously at 80 A and still stay well below 30 °C. In ordinary operating conditions with less than 30% duty cycle, the coils rest below room temperature.

The MOT coil current is supplied by an Electro-Automatik EA-PS 8080-170 power supply capable of 80 V and 170 A output. Since the settling time of the power supply is too slow (∼ 20 ms), we opt to leave the supply on constant-voltage mode and actively stabilize the current by actuating on a water-cooled Microsemi APTM50UM09FAG MOSFET installed in series with the coils – this way we achieve fast, dynamic and accurate control of the MOT current. To avoid FET breakdown we install a bank of 8 parallel TVSs to clamp the drain-source voltage to 60 V. We initially used a Powerex CM500HA-34A IGBT in place of the FET in order to use much higher clamping voltages – do not do this! It turns out that IGBTs can only handle large amounts of power if they are operated predominantly in the fully-on or fully-off state; the intermediate state in which we operate had the tendency to fuse the collector to the emitter. Our 60 V clamp is still sufficient to switch the coils from 80 to 0 A in around 1 ms. At the input stage of the control loop, the MOT current is measured using a LEM LF 205-S/SP3 current transducer; we set the coils to a particular current by feeding an analogue control voltage into a differencing stage with the current transducer output to generate an error signal to feed into the servo loop. To make sure that noise spikes from the high-power MOT coils cannot infect the experimental control electronics, an HCNR200-based analogue optical isolation stage is implemented before the FET gate-driving output. Also, on the same cooling plate as the FET, we have an H-bridge based on DM2G150SH12A IGBTs (correctly used here) to allow us to switch between Helmholtz and anti-Helmholtz current configuration so that the MOT coils can optionally be used to generate the bias/mixing field for clock spectroscopy. In practice, however, except for some early experiments we have predominantly used the auxiliary outer pair of coils for spectroscopy; the MOT coil H-bridges have largely turned out to be redundant.

During spectroscopy of the Sr clock transition, we need to be able to apply a very stable bias field of around 100 µT for ^{87}Sr and up to 2.5 mT for ^{88}Sr. This field is generated by a pair of 'mixing/bias field' coils constructed from 30 turns of 2 mm-diameter enamelled copper wire wound around the edge of the Delrin MOT coil former, yielding 125 µT/A at the atoms. These coils are driven using a bank of three parallel fan-cooled Texas Instruments OPA549 power op-amps capable of continuously supplying a total of ±20 A near the supply rails of ±12 V. The linear response of the op-amp output stage makes it a doddle to tune the control loop, especially in comparison to the highly non-linear FET-based MOT coil system; op-amps are definitely to be recommended for lower-current applications. We again install an HCNR200-based analogue opto-isolation stage between the loop filter output and the bank of power op-amps in order to avoid spikes on experimental control ground. (There is a significant mutual inductance between the MOT coils and these mixing/bias field coils, causing spikes of several tens of volts upon MOT coil switching). The control loop has a settling time of around 200 µs when switching from zero to fully on – thus we expect our field control to be limited more by eddy currents in the steel chamber ($\sim$ 1 ms time constant) rather than current in the coils. At the input stage of the control loop, we use a LEM IT 60-S flux-gate current transducer for its exceptionally high specifications and low drift. Monitoring the stability of the bias field over the course of 3 weeks using the stretched-state Zeeman splitting of the ^{87}Sr clock transition, we observe a long-term flicker floor at around 15 nT corresponding to 7 ppm of the total coil current range. This stability is most likely limited by background field fluctuations rather than coil current – the Yb^{+} lab next door measures a similar noise floor – and is easily sufficient to operate an accurate Sr clock.

A final addition, not shown in figure 2.10, is the compensation coils. There are three pairs of such coils, formed of approximately 30 turns of 0.9 mm-diameter enamelled copper wire wound around square aluminium formers. The coils are arranged around the six faces of a 60 cm cube centred on the main chamber, allowing independent compensation of the background magnetic field along all three axes. They are driven on constant current mode by a rack of six Agilent 6632B power supplies which only very rarely need to be altered.

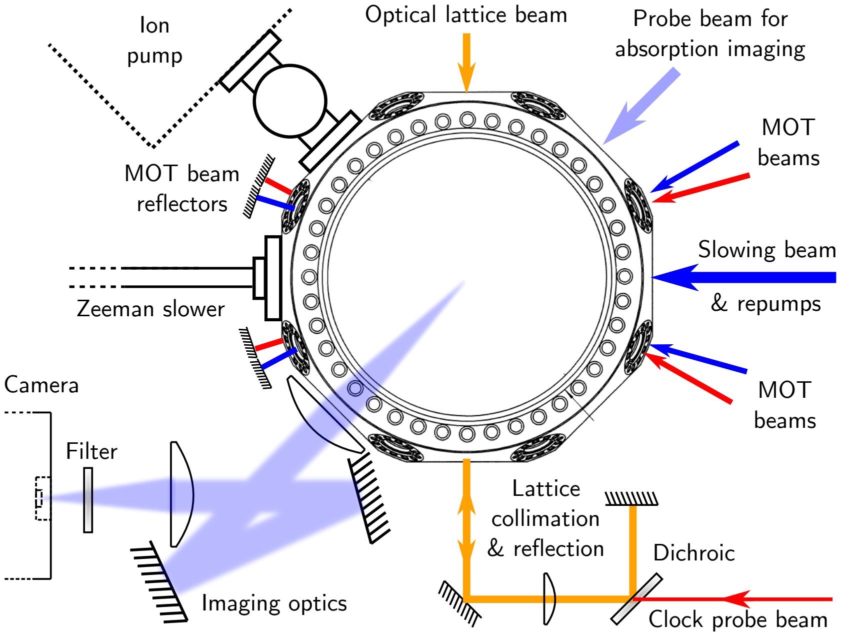

Figure 2.11: A side-on view of the science chamber. Additional ‘axial’ MOT beams propagate into the page through a gap the middle of the magnetic field coils (not shown), and an extra upward-propagating red MOT beam is sent in at a slight angle next to the clock probe.

2.2.4.2 Optical access and imaging

With so many cooling stages required to operate the Sr clock, fitting all the laser beams into the science chamber becomes quite a logistical problem. We show a simplified side-on view of the apparatus in figure 2.11, annotating many of the radial beams that are directed into the chamber. The chamber itself is sandwiched between two raised 60 × 60 cm breadboards held up by several 1.5″ diameter non-magnetic Thorlabs posts; all the optics, optomechanics, and photodetectors required to align the beams into the chamber are fixed (sometimes upside-down) either to these breadboards, to the posts, or to the optical table beneath the chamber. The purpose of the different beams will be outlined in more detail in the following sections where the individual cooling stages and laser systems are described. Here, we focus only on the chamber viewports and the imaging system.

Except for the slowing beam viewport made from c-cut sapphire, on the main chamber we exclusively use Torr Scientific zero-length fused silica viewports with MDC Vacuum Ltd AR coatings. On the MOT coil axis (into the page) we attach reducing adapters from 8″ to 2.75″ CF and install viewports with a broadband AR coating covering 461 nm to 689 nm. The coating on these viewports is also not too problematic for 813 nm with around 1.5% reflectance, making it viable to introduce a horizontal lattice beam along the MOT axis. Although the adapter to smaller viewports was chosen mostly for cost reasons, the arrangement does have the fortunate side-effects of: (1) an increased temperature homogeneity due to superior thermal conductivity of the steel adapter, and (2) a reduced viewport surface area over which potential patch charges could gather to yield a DC Stark shift.

The eight radial 1.33″ CF MOT viewports have single-layer AR coatings for the individual MOT wavelengths of 461 nm or 689 nm, while the 3 diagonal 2.75″ CF viewports, used primarily for imaging, are AR coated for 461 nm. The two 2.75″ CF viewports on the lattice axis have broadband AR coating covering 532 nm to 813 nm; these coatings were extended into the green to allow the addition of a vertical 'transport' lattice at 532 nm to move the atoms into a heatable graphite tube hanging from the top viewport − the plan was to map out the black-body radiation shift in the graphite tube as explained in [106], but such an experiment is now unlikely to be performed. For the slowing beam we avoid using a fused silica viewport because of potential erosion by the Sr atomic beam. Instead we use c-cut sapphire and apply an AR coating for 461 nm only on the outer face. This viewport is offset from the chamber by a Kimball physics close coupler, allowing access for heater tape so that the Sr on the inner face can potentially be baked off if needed. However, in our experience the Sr coating hasn't blocked the slowing beam enough to present a significant impediment to clock operation, so no serious attempt at baking has been made.

The imaging system is formed from a pair of visible-BBAR-coated achromatic doublet lenses with focal length $f = 150\,\mathrm{mm}$, and the first lens is placed a distance f from the atoms to yield a magnification of 1 onto the camera's CCD array. The CCD array on the Andor iXon Ultra camera is formed of 512 × 512 pixels each measuring $16 \times 16\,\mu\mathrm{m}$. The numerical aperture of the imaging system is limited by the viewport, with a viewing diameter of 32 mm at a distance of 110 mm from the atoms. We estimate a collection efficiency of 0.5% and a detection efficiency of 0.4% of isotropically-emitted 461 nm fluorescence photons, where this calculation uses the specified CCD quantum efficiency of slightly more than 80% at this wavelength. The CCD is routinely cooled to its minimum temperature of −80 °C to achieve the lowest possible dark count, and an electron-multiplying gain stage can optionally be used when operating at low atom number to increase the signal-to-background-noise ratio. (EM gain also introduces multiplication factor of 1.41 of electron shot noise, so it is to be avoided at higher photon counts).

The camera has various different modes of operation, of which we routinely use two. For ordinary diagnostic imaging of, e.g. time-of-flight images, we set the camera to 'Kinetics Frame Transfer' mode. This mode is appropriate whether absorption or fluorescence imaging is used, but requires enough time between pictures for the whole CCD array to be read out (typically more than 100 ms). For clock readout we can't afford to wait so long between pictures, so a different camera mode must be used. The clock readout sequence goes as follows:

- Straight after the clock interrogation pulse, a trigger pulse is sent to the camera to commence picture-taking. Then, after a 100 µs delay, the 461 nm probe light is flashed on for around 500 µs, during which the camera gathers fluorescence photons emitted by the atoms. By summing the pixel counts over a region of interest (ROI), we derive a signal g proportional to the number of atoms that were in the ground state. The ground-state atoms are blasted out of the trap by this 461 nm pulse.
- After the first picture, the repump light at 679 and 707 nm is turned on for approximately 3 ms to flush the atoms that were in the excited $^3\mathrm{P}_0$ state back to the $^1\mathrm{S}_0$ ground state.
- Next, we trigger a second picture and flash on the 461 nm probe light again for 500 µs to get a signal e proportional to the number of atoms that were excited by the clock interrogation pulse.
- Finally, in order to compensate for the background b present in both g and e we take a third picture without atoms after another 3 ms have passed.

In order to take these three pictures within $\sim$ 10 ms we have to put the camera into 'Fast Kinetics' mode, where no readout stage is implemented between the pictures. Instead, to distinguish between the three pictures the charges on the CCD are just shunted by around 30-40 pixels (taking less than 100 µs) after each picture – only once all 3 pictures have been exposed is the whole CCD array actually read out. The excitation fraction is then calculated from the pixel sums in the three ROIs as:

$$F = \frac{e - b}{(g - b) + (e - b)} \tag{2.2}$$

In practice we have found it helpful to adjust slightly the normalized detection calculation to account for the fact that the relationships between excited- and ground-state backgrounds and the pixel sum b in the final picture are not quite one-to-one.

Indeed, it seems that the effective backgrounds b_g and b_e vary significantly with ROI size, with delay between pictures, and even just with time. Thus, before setting the clock running we typically take a long series of pictures without atoms and examine the correlations between the empty g, e, and b to derive separate models fitting $b_e = m_e b + c_e$ and $b_g = m_g b + c_g$ for some constants m_e, c_e, m_g, c_e. Thankfully, however, accurate calibration of the backgrounds isn't all that critical for an accurate clock – we tend to use methods of interrogation which are intrinsically immune to background offsets (see section 3.2.1).

2.3 Cooling sequence and lasers

The basic cooling sequence for Sr goes as follows: we initially capture atoms into a blue MOT at 461 nm, then we transfer into a red MOT at 689 nm for further cooling, and finally we load cold atoms from the red MOT into the lattice trap for clock interrogation. The whole sequence is summarized with reference to the Sr term diagram in figure 2.12, and we use the rest of this section to elaborate on our experimental implementation of each stage.

2.3.1 The Zeeman slower and blue MOT

The first stages of generating, slowing, and capturing Sr using the ${}^1\mathrm{S}_0 \to {}^1\mathrm{P}_1$ transition at 461 nm had already been demonstrated before I joined the project, and in large part the apparatus remains the same as described in previous theses [46, 106]. Perhaps the most major improvement is the addition of new MOT coils to enable reliable continuous operation, but there has also been some significant rearrangement of the 461 nm optics and electronics for beam distribution and delivery.

Initially, the Sr beam exits from a 1 mm-diameter single-channel nozzle at the end of an oven crucible which, immediately after loading, holds around 4 g of 99.9%-purity dendritic strontium from Sigma Aldrich. The stainless-steel oven is heated in vacuum by a heater wire to a typical operating temperature of 600 °C, aided by the presence of two layers of heat shields that allow operation with a modest input power of 28 W. The wide single-channel nozzle is comparatively wasteful of atoms, so the oven lifetime between refills is only about a couple of years. The atomic beam exits the oven with a distribution of velocities centred around a peak of $\sim 550\,\mathrm{m\,s^{-1}}$ and then passes into the Zeeman slower tube through a 1 mm-diameter aperture 10 cm downstream of the oven nozzle. In the transverse-field spin-flip permanent-magnet Zeeman slower, proposed in [206] and characterized in [107, 108], we address the portion of atoms

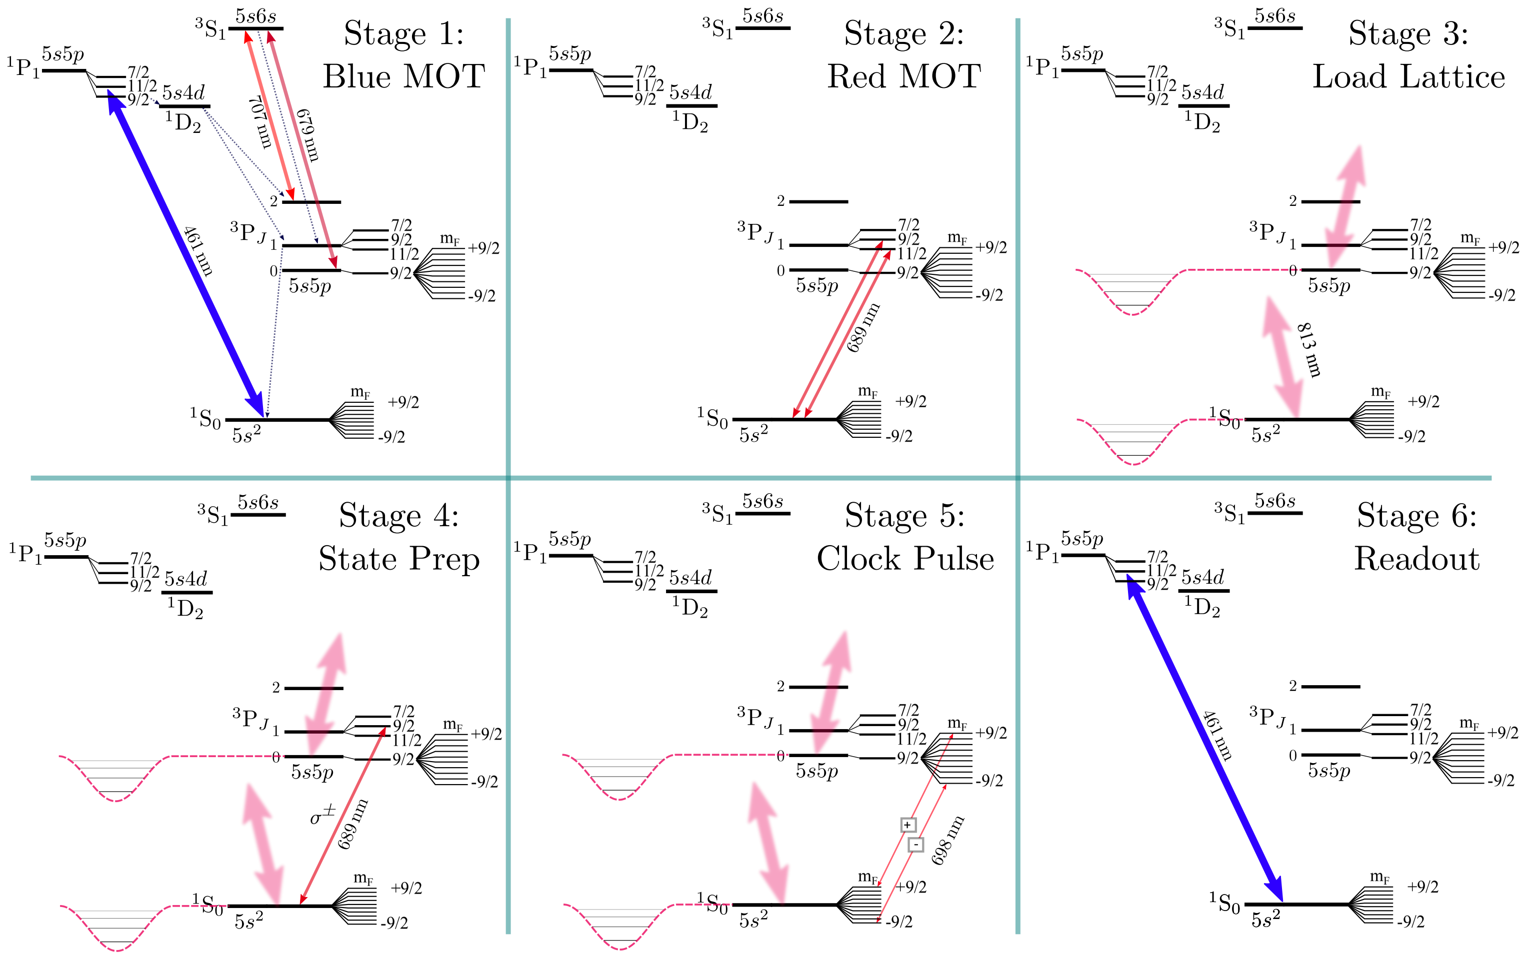

Figure 2.12: Cooling sequence and clock interrogation stages.

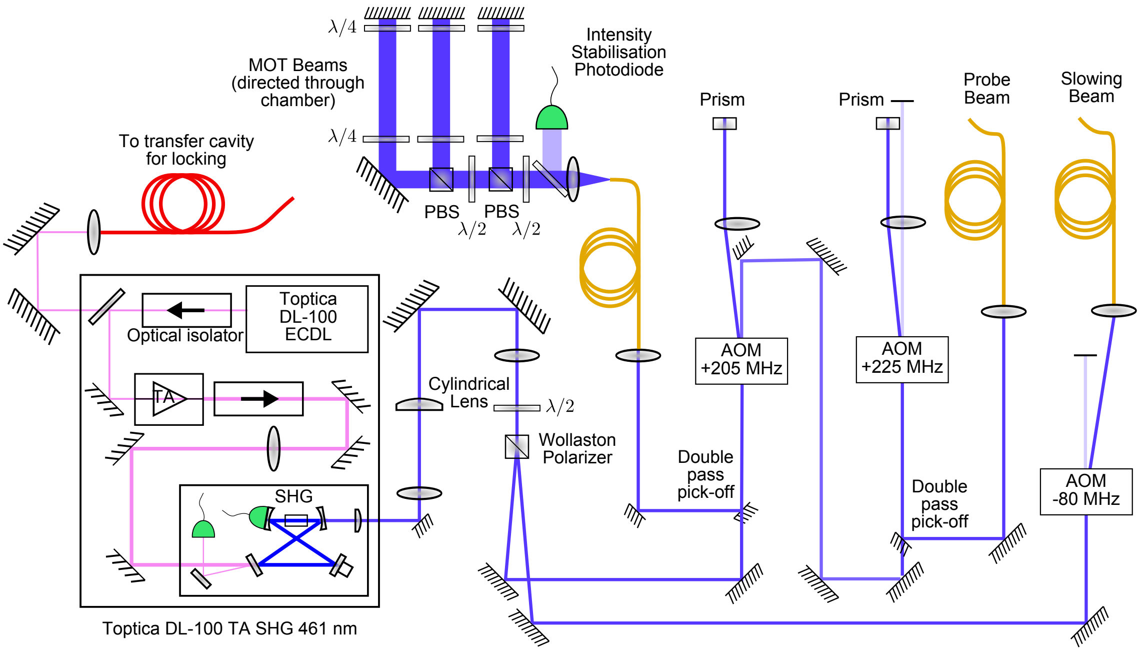

Figure 2.13: Our optical system for generation and distribution of 461 nm light. In addition, we install shutters before the fibres for complete extinction of stray 461 nm light during clock spectroscopy. The diagram is adapted from [106].

exiting the oven at $\lesssim 500\,\mathrm{m\,s^{-1}}$ using a fibre-delivered slowing beam with up to 60 mW power at a detuning of around −540 MHz from the 461 nm transition; this yields a flux of atoms exiting the slower at around $35\,\mathrm{m\,s^{-1}}$ into the main science chamber.

After the atoms enter the science chamber, we implement a sequence of steps to capture and cool them in preparation for clock interrogation. The first capture stage is based on a broad-line magneto-optical trap (MOT), the standard technique of laser cooling originally demonstrated in [217] and described in many atomic physics textbooks ever since. During the capture stage, we operate a 'blue' MOT at 461 nm with up to 6 mW of power in each of the three incident beams at a detuning of around −40 MHz and with a coil-axis magnetic field gradient of 4 mT/cm. The beam waist is $w_0 \approx 7\,\mathrm{mm}$, corresponding to a peak saturation of up to $I/I_{\mathrm{sat}} = 0.2$ per beam ($I_{\mathrm{sat}} = 42\,\mathrm{mW/cm^2}$). The beam geometry is not quite orthogonal (see figure 2.11), so the cloud of atoms tends to be much more extended vertically than horizontally. After loading enough atoms into the MOT (this typically takes 60 − 600 ms) we implement a brief 10 ms blue cooling stage. In this stage the slowing beam is turned off, and the MOT beams are jumped to −25 MHz detuning and linearly ramped to zero intensity. This gives a chance for the atoms to move into the centre of the quadrupole magnetic field – during loading, they are typically pushed by around 3-4 mm towards the oven by the slowing beam – and it gets the atoms as cold as possible for efficient transfer into the second stage of cooling.

The 461 nm laser system is depicted in figure 2.13. The light is generated from a master oscillator plus power amplifier (MOPA) system at 922 nm which is frequency doubled by a $KNbO_3$ crystal held at 39 °C inside a bow-tie cavity. This Toptica laser system is still capable of generating 350 mW at 461 nm even after eight years of near-continuous use, though the doubling crystal has occasionally had to be moved to restore maximum power, and the output of the 922 nm master laser diode has steadily declined to $\sim 80\%$ of its original specifications despite now being operated 10% above its nominal maximum current. In contrast with notoriously unreliable 813 nm tapered amplifiers, the 922 nm TA has shown no sign of deterioration at all. The frequency of the 922 nm master laser is indirectly stabilized to the clock laser via a PDH lock to the transfer cavity (see section 2.2.1), with modulation applied directly to the diode current rather than an EOM – this leads to worsened offsets from residual amplitude modulation, but the lock accuracy is still more than sufficient to address the 30 MHz-linewidth atomic transition. After exiting the laser the 461 nm light is sent through a telescope with a cylindrical lens for circularization and collimation, and is then split between the slowing beam and the MOT/probe beams on an optically-contacted quartz Wollaston polarizer from Leysop Ltd. This polarizer is chosen for its power-handling capability: the high peak intensity of the 461 nm beam kept burning holes in the cemented polarizing beamsplitters from Thorlabs. The AOMs are all IntraAction ATM series TeO_2 crystal with 1 mm active apertures, driven by Minicircuits ZOS series VCOs amplified by Minicicuits ZHL-2010 RF amplifiers, typically achieving 90% maximum single-pass diffraction efficiency. We intensity stabilize the probe and MOT beams, which are both used to create short imaging pulses, by actuating on a Minicircuits ZFAS-2000 bi-phase/variable attenuator placed before the RF amplifier, chosen for its high 500 kHz modulation bandwidth in order to achieve minimum settling times. (We previously found ourselves limited by the ~ 20 kHz bandwidth of the ZX73-2500 variable attenuators used as standard elsewhere in the lab).

We operate the blue MOTs for ^{87}Sr and ^{88}Sr at different frequencies to account for the isotope shift. For ^{87}Sr the $^1S_0$ (F=9/2) $\rightarrow$ $^1P_1$ (F=11/2) and $^1S_0$ (F=9/2) $\rightarrow$ $^1P_1$ (F=9/2) transitions are shifted by -55.2 and -72.2 MHz respectively from the $^1S_0 \rightarrow {}^1P_1$ transition in ^{88}Sr. The presence of the extra $9/2 \rightarrow 9/2$ transition slightly conflicts with the desire to operate an ^{87}Sr MOT on the main $9/2 \rightarrow 11/2$ transition, but the conflict isn't prohibitive: at typical -40 MHz detunings from the main MOT

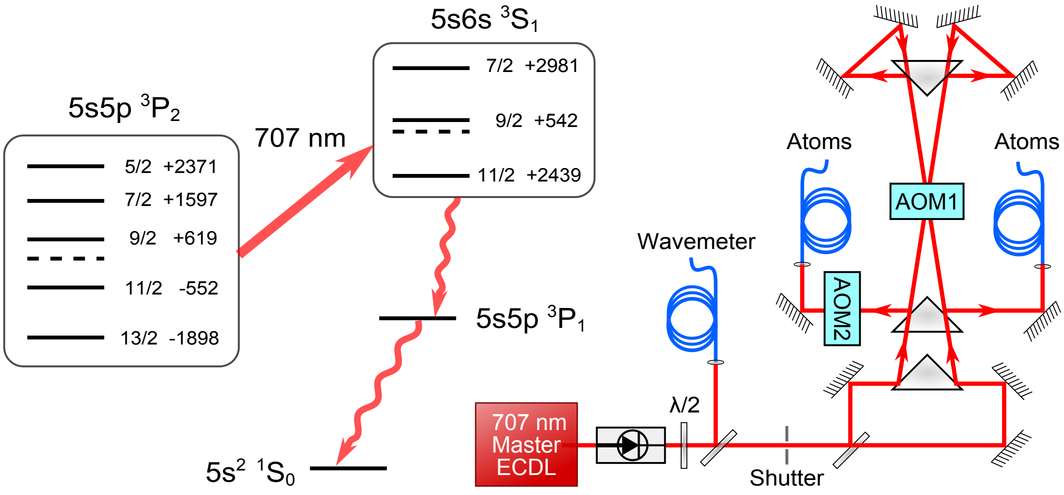

Figure 2.14: Our optical system for generation and distribution of 707 nm repump light covering all hyperfine levels of ^{87}Sr. See text for explanation of the AOM setup. Hyperfine splittings are taken from [243].

transition, the beams still form an effective molasses on the $9/2 \rightarrow 9/2$ transition at -23 MHz detuning. Since there is no bridging AOM between the 922 nm master and the transfer cavity (see figure 2.2), we must tune the cavity itself in order to select between blue MOTs with ^{88}Sr or ^{87}Sr. Thus, when changing from ^{88}Sr to ^{87}Sr we adjust the 698 nm double-passed offset AOM by -19 MHz to effect a -56 MHz shift at 461 nm, matching the isotope shift of the main MOT transition. We find that the cavity adjustment is not all that sensitive within around ± 10 MHz, as long as the slack is taken up by the MOT and probe AOMs, implying we can achieve similar MOT loading rates over a large range of slowing beam detunings.

2.3.1.1 Repumping

From the $^1P_1$ excited state in the blue MOT there is a significant branching ratio into the the metastable $^3P_{0,2}$ states of about 1 in 50000, limiting the MOT lifetime to ~ 10 ms for typical $I/I_{\text{sat}} \sim 0.1$ in each beam. To extend this lifetime and improve the blue MOT atom number, repumps are applied throughout the loading and cooling stage. There are several possibilities for efficient repumping [183, 246], but we follow the most common route by driving the 679 and 707 nm transitions to 5s6s $^3S_1$ [148, 292], recycling the atoms to ground via the $^3P_1$ manifold within around 21 µs. The 13 ns lifetime of the $^3S_1$ manifold yields a natural linewidth of 12 MHz for these repumping transitions, meaning that frequency stabilization to a HighFinesse WSU series wavemeter with nominal 5 MHz accuracy is quite sufficient. We stabilise the repump frequencies using a LabView software-based PI control loop actuating on the ECDL piezo. However, when this lock fails (usually due to laser mode-hops) the software is capable of jumping the laser diode current until the piezo lock can successfully be recaptured – this automatic relocking is essential to allow the system to operate unsupervised over long periods of more than a day. We note that the repump laser diodes are quite fussy, especially the AR-coated 707 nm diode which lases naturally at around 690 nm, requiring very careful alignment of the grating feedback and excellent mechanical stability – our system is close to good enough, but replacement with short interference-filter-based ECDLs is underway for even better long-term reliability (only 10 years behind the SYRTE goup [14]).

In the ^{87}Sr isotope, the presence of nuclear spin complicates the repumping process by introducing 5 separate hyperfine manifolds within the metastable $^3P_2$ fine structure manifold, spread in energy over more than 4 GHz. A single-frequency laser simply isn't enough to address all these transitions, leading to rather poor MOT lifetimes and equilibrium ^{87}Sr MOT numbers. One approach to efficient repumping is simply to modulate the laser by a few GHz on the piezo, as successfully applied in [43]. However, our ECDLs have comparatively bulky grating mounts which can't be strongly modulated at a rates much above 100 Hz: when we tried the laser modulation approach we achieved a disappointing $\sim 3-4$ atom number enhancement at best. Instead, we implemented the AOM setup depicted in figure 2.14 to generate sidebands directly addressing transitions from all 5 hyperfine manifolds in $^3P_2$. AOM1 is driven at 475 MHz and two incident beams are aligned to the positive and negative Bragg angles and double-passed to provide individual frequency components at -950, -475, 0, $+475$ and $+950$ MHz (corresponding to the -2, -1 , 0, $+1$ and $+2$ diffraction orders of the AOM). The -2, -1 and 0 orders are picked off and frequency-shifted by -225 MHz by AOM2 before being delivered to the atoms, while the $+1$ and $+2$ orders are both sent directly via another fibre. When the 707 nm master is frequency-locked via the wavemeter, the resulting frequency components at -1175, -700, -225, $+475$ and $+950$ MHz are all within 20 MHz of the $F = 7/2 \rightarrow 9/2$, $13/2 \rightarrow 11/2$, $9/2 \rightarrow 9/2$, $5/2 \rightarrow 7/2$ and $11/2 \rightarrow 9/2$ transitions respectively, ensuring that the atoms have nowhere left to hide. With between 350 µW and 1 mW in each mode at the atoms, and applying some current-modulation of the 707 nm master laser for extra frequency broadening, the ^{87}Sr blue MOT lifetime remains reliably above 1 s. Due to an underwhelming slow flux of ^{87}Sr, we have found this MOT lifetime enhancement indispensable in maintaining decent lattice-trapped atom numbers during long-term clock operation.

As a further technical note, it proved important in our system to deliver the repumping beams along the slowing-beam axis: we find that the blue MOT is pushed during the loading stage by the slowing beam passing through it, so the MOT jumps by $\sim 3-4$ mm along the slowing beam axis at the end of the loading stage. If the repumpers are delivered from the side, they have to be much bigger (and therefore less intense) to cover both blue MOT regions – co-propagation with the slowing beam prevents this constraint.

2.3.2 The red MOT

The spin-forbidden $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_1$ transition at 689 nm in strontium provides a rare opportunity for efficient laser cooling to ultracold temperatures, with a narrow 7.5 kHz linewidth, yielding an extraordinarily low Doppler cooling limit of 180 nK. In this unusual regime of Doppler cooling, the temperature is limited by the single-photon recoil energy $E_r = \hbar^2 k^2/2m = k_B \times 230\,\mathrm{nK}$ [51], though an intriguing possibility remains of sub-recoil cooling with modulated light [274]. Laser cooling to sub-µK temperatures in a 689 nm 'red' MOT was first pioneered in Tokyo with ^{88}Sr [119, 129] and later with the more difficult fermionic ^{87}Sr isotope [189]. More recently, the transition has been exploited to prepare quantum degenerate gases of all stable Sr isotopes [245, 247] in a cooling sequence which can be as short as 2 s in the most propitious case of ^{84}Sr – the world's fastest BEC [243, 244]. The dynamics of the red MOT have been studied in detail in [163] and are thoroughly explained in various doctoral theses such as [43, 164].

In practice, the narrow linewidth of the 2$^{\mathrm{nd}}$-stage cooling transition is not universally a blessing. The first problems emerge in the transfer stage from the blue MOT: the atoms left over once the blue MOT beams are switched off are typically at a temperature of ~ 2 mK, with a 1/e spread in velocity of $v_{\mathrm{th}} \approx 0.7\,\mathrm{m\,s^{-1}}$ leading to Doppler broadening at 689 nm of around 1 MHz, much wider than the natural 7.5 kHz linewidth of the transition we're aiming to use. The established solution to this problem is to frequency-modulate the 689 nm light in an initial 'broadband' red MOT [129] so that some red light is resonant with all but the fastest atoms. In our implementation we modulate the MOT beams at a 40 kHz rate covering an optical spectrum of 4 MHz peak-to-peak, with a detuning of approximately −100 kHz in the nearest sideband. With all parameters optimized, this modulation helps us achieve roughly 40% transfer efficiency from the blue into the red MOT. (We also modulate the 'stirring' beam at the same rate but over a 2 MHz optical spectrum – this beam is needed for ^{87}Sr to drive the less magnetically sensitive $\mathrm{F} = 9/2 \rightarrow \mathrm{F} = 9/2$ transition which randomises the m_F magnetic sublevel, significantly reducing atom loss rates in the ^{87}Sr red MOT [189]). A further technical problem, arising in the blue-red MOT transfer stage, is due to the magnetic field gradient: in order to operate an efficient blue MOT we must use high gradients of around 4 mT/cm, which wildly shifts the red MOT transition with a sensitivity of $\mu = 21\,\mathrm{MHz/mT}$. For this reason we need to switch the field gradient to a much smaller value of around 0.2 mT/cm for the broadband red MOT stage, and in order to stop the atoms from escaping before the red MOT can take hold we need this field to switch fully within ~ 2 ms of the end of the blue MOT (hence the need for our fast MOT coil driver described in section 2.2.4.1).

The atoms captured in the broadband red MOT typically reach a temperature of 10-20 µK within 80 ms, compressing into a cloud of around 1 mm diameter. To get even colder and denser samples, we then implement a single-frequency red MOT stage, turning the 689 nm modulation off and considerably reducing the beam intensity. In this stage another complication arises from the narrow transition linewidth: even if we fully saturate the $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_1$ transition it can only scatter photons at a rate of $23\,000\,\mathrm{s}^{-1}$ $(\Gamma/2)$, implying a maximum acceleration of $155\,\mathrm{m\,s}^{-1}$ when pushed in one direction by a single beam. With the scattering force being this weak, the atomic acceleration under gravity begins to play a substantial role in the MOT dynamics. In our case, the dominance of gravity is amplified by a rather unwise choice of beam geometry – as shown in figure 2.11, the MOT beams are not directed along three orthogonal axes; instead, the radial MOT beams are sent in at an angle of 22.5° to horizontal, reducing our maximum upward acceleration to just a factor of 8 above gravity. If we further take into account saturation by the axial MOT beams as well as the unavoidable absorption of the occasional downward-propagating photon, we find that the atoms can only barely be supported under their own weight.

To solve the gravity problem we introduced an extra upward-propagating red MOT beam (the ‘up’ beam). Even though we couldn’t retro-reflect this beam due to the narrow graphite tube hanging off the top viewport [106], we have found the up beam immensely helpful: in ^{88}Sr it lets us create long-lived single-frequency red MOTs below 1 µK, while in the less efficient ^{87}Sr red MOT it is the only way we can make cold single-frequency MOTs with a lifetime of any more than 20 ms.

Since there is significant literature already available on red MOTs in strontium, we haven’t gone through any rigorous characterization of 2$^{\mathrm{nd}}$-stage cooling. For our purposes, it is sufficient just to optimize directly the experimental control parameters (e.g. intensity setpoint voltages, detuning control voltages, timings, and beam

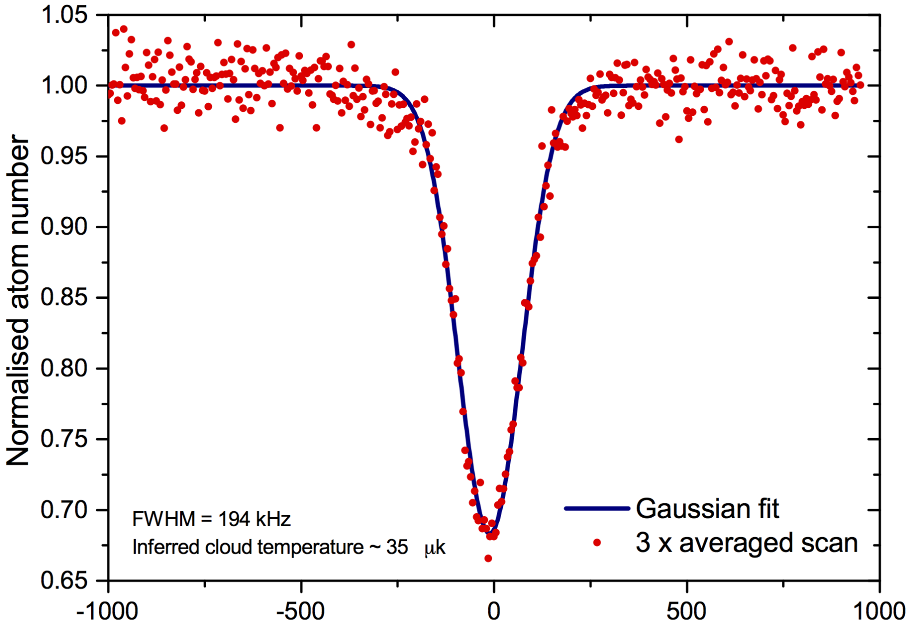

Figure 2.15: Our first sighting of the clock transition: atom loss in the ^{87}Sr red MOT. A 35 µK atom temperature is inferred from the width of the Doppler-broadened line.

alignments) for loading of cold Sr into the optical lattice trap. The small amount of nice-looking data we do have is shown in figures 2.16 and 2.17 showing ^{88}Sr red MOTs, and figure 2.15 showing early clock spectroscopy in the ^{87}Sr MOT – notable only because it was the first time we had proof that the clock laser actually did anything.

Soon after we made our first ^{88}Sr red MOT, we characterised its temperature using the time-of-flight images in figure 2.16. From the rate of expansion of the cloud we calculate horizontal and vertical temperatures of 1.1 µK and 1.8 µK respectively

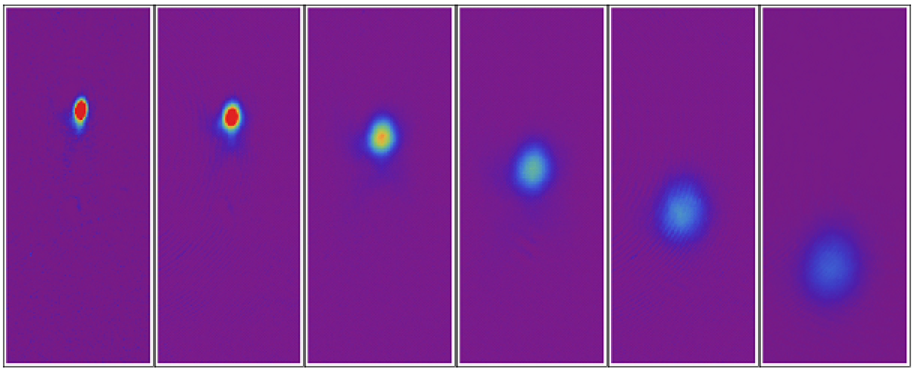

Figure 2.16: Red MOT time-of-flight images: 1 ms, 5 ms, 10 ms, 15 ms, 20 ms, 25 ms.

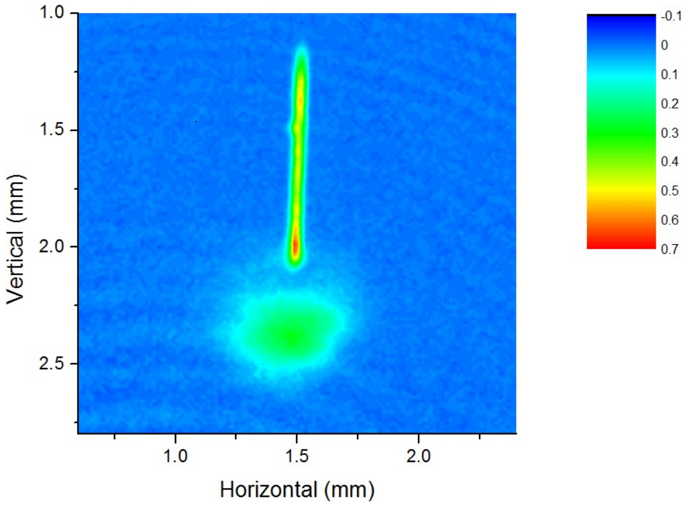

Figure 2.17: Absorption image in units of optical depth from a ‘streak loaded’ lattice, showing the remnants of a red MOT below the string of trapped atoms. The image is taken at a 45 degree angle from the vertical axis, so the implied loading region is around 1.4 mm in length covering around 3500 lattice sites.

with up to 1.5×10^7 atoms. These data were taken prior to the installation of the up beam (hence the horizontally-squeezed profile of the cloud) using an intensity of $I/I_{\mathrm{sat}} = 30$ per MOT beam, where $I_{\mathrm{sat}} = 3\,\mu\mathrm{W/cm}^2$. Later, using the up beam, we were able to ramp to much lower intensities to achieve $\sim 700\,\mathrm{nK}$ temperatures in long-lived lenticular red MOTs more closely resembling those in [163].

One novel development in our system was to construct a ‘streak loading’ sequence for ^{88}Sr, with results shown in figure 2.17. In this sequence we linearly ramp the detuning of the red MOT beams in the final stage by up to $-1\,\mathrm{MHz}$ for $50 - 300\,\mathrm{ms}$. During the ramp, the atoms sag under gravity into a position where the magnetic field shifts the σ^+ transition into resonance with the upward-propagating beams. If the detuning ramp is slow enough then the atoms fall in a controlled manner, loading a long string of lattice sites in their wake. Note that we observe streak loading only when the red MOT beam intensity is $\lesssim 10 I_{\mathrm{sat}}$: we theorise that higher-intensity beams flush out the atoms that were initially trapped in the lattice, so at high intensity we observe remaining lattice-trapped atoms only at the position where the red MOT ramp ends. With careful tuning of the MOT intensity, however, it is possible to load reliably over > 1000 lattice sites, significantly reducing the atomic density and therefore preventing collisional shifts and dephasing of the clock transition (see section 4.2.1).

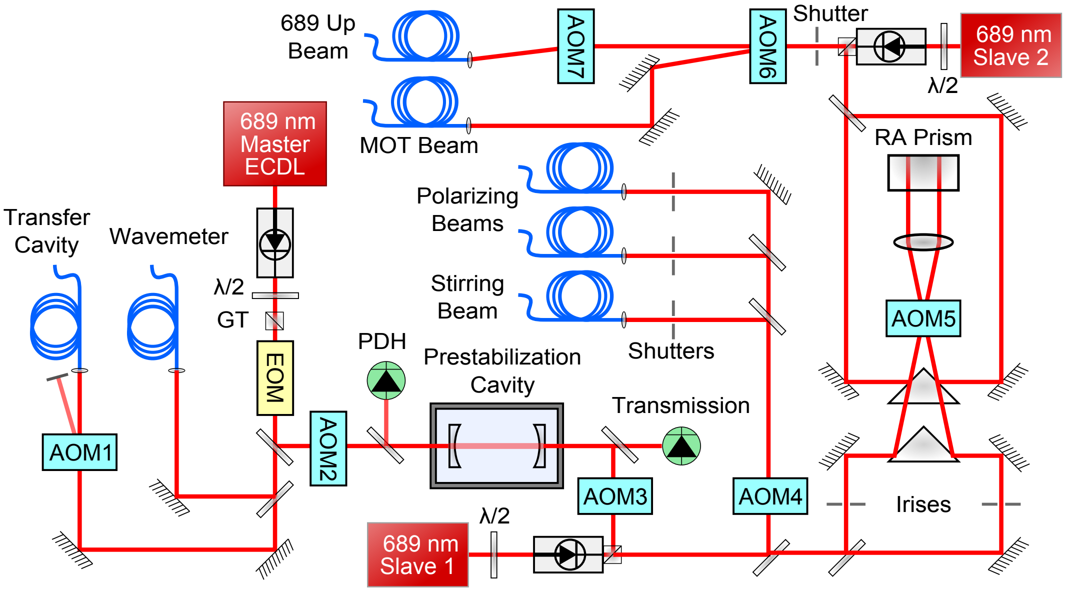

Figure 2.18: Our optical system for generation and distribution of 689 nm light. AOM1: RAM servo. AOM2&3: Frequency bridge between cavity resonances (shared RF source, double-passed). AOM4: Intensity stabilization, bridge to stir transition, and FM (double-passed). AOM5: Isotope-bridge and FM (select isotope using irises). AOM6&7: Intensity stabilization and bridge to MOT transition (shared RF source).

2.3.2.1 689 nm laser system

We depict the 689 nm laser system in figure 2.18. The master laser is a Littrow-configuration extended-cavity diode laser (ECDL) designed by Ian Hill [106], similar to the clock ECDL depicted later in section 3.5. The master laser diode is a 'virtual point source laser' VPSL-0690-035-X-5-A from Blue Sky Research, containing an integrated microlens to produce a convenient circular output. The laser linewidth is narrowed using a 2 MHz-bandwidth PDH lock to the prestabilization cavity, actuating on the ECDL current (fast) and piezo (slow). A much slower 50 mHz-bandwidth PDH lock then actuates on AOM2 to steer the master (and the 1^{st} slave) towards the transfer cavity resonance. As described in detail in section 2.2.1.2 we implement a RAM servo to prevent drifts in the offset of the transfer cavity lock. The pre-stabilization cavity uses a 10 cm ULE spacer and has a finesse of 65000 at 689 nm, resulting in a 23 kHz cavity linewidth; this narrow feature makes it easy to narrow to kHz laser linewidth and it also helps to filter out the servo sidebands on the transmission which could otherwise affect the atoms during the red MOT cooling stage. The pre-stabilization cavity is almost good enough as a reference by itself, but unfortunately it is operated around 22 °C above the zero in its coefficient of thermal expansion, resulting in frequency swings at the scale of 100 kHz over an hour; thus we need to introduce the transfer cavity lock to achieve the required sub-10 kHz long-term stability.

We use the 70 µW beam transmitted through the pre-stabilization cavity to seed two injection-locked slave diode lasers (both are Opnext HL6750MG diodes with up to 50 mW output). Slave 1 is the source of the light addressing the $F = 9/2 \rightarrow 9/2$ 'stir' transition in ^{87}Sr. This transition is used in two different stages of ^{87}Sr atom preparation: (1) spin-randomization in the red MOT to improve MOT lifetime (as discussed above), and (2) spin-polarization of the atoms into a single ground-state magnetic sub-level $M_F = \pm 9/2$ before clock interrogation (see figure 2.19). In both applications, the stirring beams are intensity stabilized by actuation on AOM4; this AOM is also used to modulate the frequency of the light by 2 MHz during the broadband MOT and the spin-polarizing stages. We select between the tasks of spin-randomization or spin-polarization using the shutters placed before each separate fibre. Most of the stirring light is diverted for the spin-randomization beam so that we can use up to 8 mW during the broadband MOT. The remaining power is easily sufficient for spin-polarization, which needs less that $\sim 300\,\mu$W in each beam.

More optics, not shown on the diagram, are required to handle the stir-transition beams at the remote end of the fibres (around the science-chamber). The spin-randomisation beam is combined with the main red MOT beam on a polarising beam-splitter, before getting split three ways to form a 3D optical molasses. Meanwhile, the spin-polarisation beams are manipulated to ensure good, stable σ^+/σ^- polarisations incident on the atoms – at the fibre outputs, we collimate to a beam radius of ~ 1.5 mm, we recombine the two beams in free space on a polarising beamsplitter, we pick them off onto the same photodetector for intensity stabilisation, and we finally send them through a quarter-wave plate down the quantisation axis onto the atoms.

Slave 2 is the source of light for the main F $\rightarrow F+1$ MOT transition in both ^{87}Sr and ^{88}Sr. We select between the two isotopes using the irises into AOM5: This AOM is operated at ~ 310 MHz so that the -2 and $+2$ diffraction orders are separated by the MOT transition isotope shift of 1241.452 MHz. The isotope-bridge AOM is also used to modulate the frequency of the light injecting slave 2 (we checked via a beat with the master ECDL that the slave easily follows the modulation). AOM6&7 are for intensity stabilization of the MOT beams, used in single pass to maximize the power available for the broadband MOT – we get up to 14 mW delivered to the science chamber.

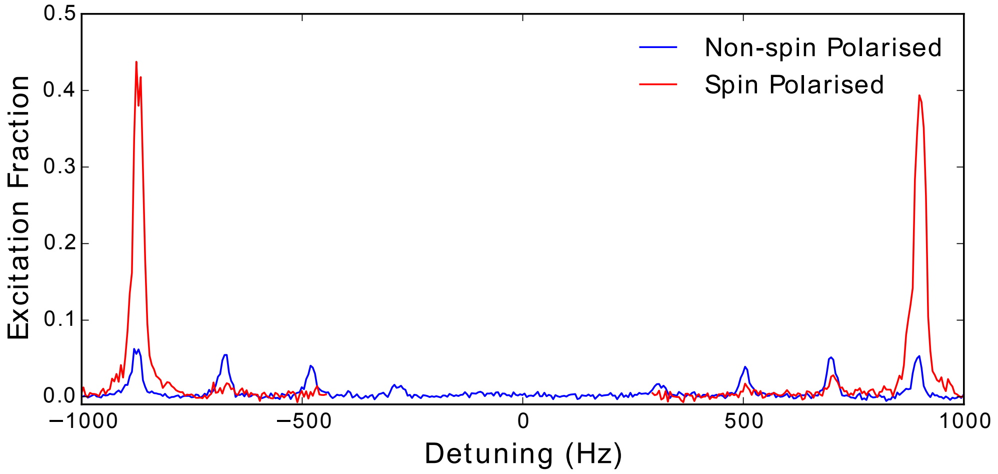

Figure 2.19: Scans over the Zeeman sublevels of the ^{87}Sr clock transition using π-polarised clock light, showing the effect of a spin-polarisation stage at 689 nm before the clock interrogation pulse. Note that there are two separate red traces on the plot, corresponding to polarisation into the $M_F = +9/2$ or the $M_F = -9/2$ ground state.

2.3.3 The optical lattice trap

Throughout the red MOT, we continuously shine an optical lattice trap onto the atoms. By the end of the MOT stage, the Sr cloud is cold and dense enough that a significant fraction of the atoms load automatically into the trap. Once loaded, the atoms can be kept for anywhere up to several seconds while we carry out clock spectroscopy.

In this section we briefly outline the operating principles of the optical lattice trap, we explain its purpose in facilitating accurate clock spectroscopy, and we describe its implementation in our experiment.

2.3.3.1 Optical lattice basics

When we shine far off-resonant laser light at the strontium atoms in the ground state, the light's electric field $\mathbf{E}(\mathbf{x},t) = \mathbf{E}(\mathbf{x}) \cos \omega t$ causes a position-dependent Stark shift $U(x)$:

$$U(\mathbf{x}) = -\alpha_g(\omega) \left(\frac{|\mathbf{E}(\mathbf{x})|}{2} \right)^2 \tag{2.3}$$

where $\alpha_g(\omega)$ is the scalar polarizability of the ground state at the light's frequency ω (for now, we ignore the much smaller vector and tensor terms – see section 4.1). The Stark shift exerts a dipole force $\mathbf{F}_{\text{dip}} = -\nabla U(\mathbf{x})$ which, assuming we use a red-detuned laser frequency such that $\alpha_g(\omega) > 0$, confines the atoms to a region around the local maximum beam intensity. A focused, off-resonant laser beam is thus appropriately called a 'dipole trap' [94].

An optical lattice is a specific type of dipole trap in which an optical standing wave is used to make a series of closely spaced wells in the potential $U(\mathbf{x})$ as depicted in figure 2.20. In the figure, we show the most frequently-used geometry for optical lattice clocks – the one-dimensional lattice – which can be formed using a single Gaussian beam reflected back onto itself. To a good approximation, the 1D lattice potential can be written:

$$U(r,z) \approx -U_0 e^{-\frac{2r^2}{w_0^2}} \cos^2 k_L z \tag{2.4}$$

where r and z are the radial and axial directions with respect to the laser beam, $k_L = 2\pi/\lambda_L$ is the lattice k-vector, and U_0 is the trap depth given by:

$$U_0 = \alpha_g(\omega) \frac{4P}{\epsilon_0 c \pi w_0^2} \tag{2.5}$$

where P is the power in each of the counter-propagating beams and w_0 is the beam waist. For simplicity, the Guoy phase, radius of curvature, and expanding waist of the Gaussian beam have all been ignored – the atoms are typically loaded near the beam waist where these parameters have no significant effect.

We observe that the atoms are attracted to a series of pancakes ('lattice sites') along the z axis, corresponding to antinodes in the standing wave spaced at half the laser wavelength $\lambda_L/2$. In the approximation that the atoms are tightly confined to the middle of each site, the sites can be modelled as a harmonic potential with characteristic trap frequencies given by:

$$\omega_z = 2\pi\sqrt{\frac{2U_0}{m\lambda_L^2}} \qquad\qquad \omega_r = 2\pi\sqrt{\frac{U_0}{m(\pi w_0)^2}} \tag{2.6}$$

where m is the mass of a single Sr atom, and ω_z and $\omega_r (= \omega_x = \omega_y)$ denote the axial and radial trap frequencies respectively.

For reasons explored in section 2.3.3.3, we use a lattice at $\lambda_L = 813\,\mathrm{nm}$ for clock spectroscopy, leading numerically to an axial frequency of:

$$\omega_z/2\pi = \sqrt{\frac{U_0}{35\,\mu\mathrm{K}}} \times 100\,\mathrm{kHz} \tag{2.7}$$

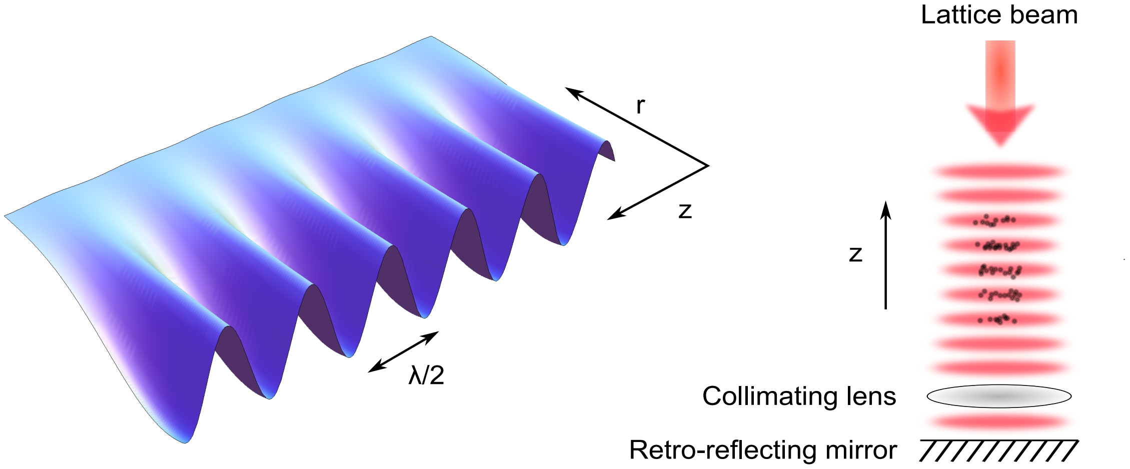

Figure 2.20: Illustration of the 1D optical lattice trap and its potential. Note that the sketches are not to scale: in reality the lattice sites are much more squished together at $\lambda_L/2 = 407\,\text{nm}$ separation, the waist is relatively much wider at $45\,\mu\text{m}$, and the overall beam is around 50 cm long from the atoms to the retro-reflecting optics.

so that over the range of typical trap depths $U_0 \approx 1 - 150\,\mu\text{K}$ we can achieve axial frequencies $\omega_z/2\pi \approx 15 - 200\,\text{kHz}$.

A typical beam waist is around $w_0 \approx 30 - 150\,\mu\text{m}$, much larger than the lattice-site spacing of 407 nm, meaning that the atoms are significantly less well confined in the radial direction than along the z axis. Crunching the numbers, the atoms therefore slosh around much more slowly with a radial trap frequency:

$$\omega_r/2\pi = \sqrt{\frac{U_0}{10.6\,\mu\text{K}}} \times \frac{100\,\mu\text{m}}{w_0} \times 100\,\text{Hz} \tag{2.8}$$

Note that these harmonic-oscillator equations only hold in the approximation of tight atomic confinement into a single lattice site. For the full lattice potential given in equation 2.4, several complications emerge. For example, (1) the radial potential is in fact Gaussian rather than quadratic, causing an energy spacing that diminishes with quantum number n_r, (2) since the lattice beam is less intense as we move away from its axis, we see a reduced axial trap frequency with increased n_r (visible in the sideband scans in figure 2.22) [32], and (3) the periodicity of the lattice encourages tunnelling between the lattice sites, with axial eigenstates best described either by Bloch or by Wannier-Stark wavefunctions depending on lattice orientation with respect to gravity [153]. The harmonic approximation holds best when the atoms are cold, when the trap is deep, and when the lattice sites possess non-degenerate energy (achieved e.g. using a lattice oriented along gravity).

2.3.3.2 Why a lattice? The Lamb-Dicke regime

Dipole trapping of the atoms during clock interrogation is essential for extending probe times out to 100 ms and beyond – without a trap, the cloud of atoms would quickly expand and fall out of the clock interrogation region. The dipole trap therefore lets us reach a much higher transition Q, increasing the frequency sensitivity of the atomic response and ultimately yielding a more stable clock (see chapter 3).

However, tight atomic confinement in an optical lattice also has another important advantage – it freezes out the Doppler and recoil shifts. To explain this effect, which is also very prominent in trapped-ion systems, we write out the atomic Hamiltonian in the harmonic approximation:

$$\begin{aligned}\hat{H}_0 &= \hbar\omega_0 \left|e\right\rangle\left\langle e\right| + \sum_{i=1}^{3}\left(\frac{\hat{p_i}^2}{2m} + \frac{1}{2}m\omega_i^2\hat{x_i}^2\right) &(2.9)\\ &= \hbar\omega_0 \left|e\right\rangle\left\langle e\right| + \sum_{i=1}^{3}\hbar\omega_i\left(\hat{a}_i^\dagger\hat{a}_i + \frac{1}{2}\right) &(2.10)\end{aligned}$$

where the ω_i are the lattice trap's angular frequencies along the three spatial dimensions (numerically calculated in section 2.3.3.1 above), and the $\hat{a}_i^\dagger$ and $\hat{a}_i$ are the well-known raising and lowering operators [59]:

$$\begin{aligned}\hat{x}_i &= x_i\left(\hat{a}_i^\dagger + \hat{a}_i\right) \qquad \text{with} \qquad x_i = \sqrt{\frac{\hbar}{2m\omega_i}} &(2.11)\\ \hat{p}_i &= ip_i\left(\hat{a}_i - \hat{a}_i^\dagger\right) \qquad \text{with} \qquad p_i = \sqrt{\frac{m\omega_i\hbar}{2}} &(2.12)\end{aligned}$$

where x_i, p_i (without the hats) describe the characteristic length and momentum scales along each dimension.

The energy eigenstates in the harmonic oscillator are the number states $\left|\mathbf{n}\right\rangle = \left|n_x, n_y, n_z\right\rangle$, which follow

$$\hat{a}_i^\dagger\hat{a}_i\left|\mathbf{n}\right\rangle = n_i\left|\mathbf{n}\right\rangle \tag{2.13}$$

and thus possess energy:

$$E_{\mathbf{n}} = \sum_{i=1}^{3} \hbar\omega_i \left(n_i + \frac{1}{2}\right) \tag{2.14}$$

When we introduce a clock probe beam with an electric field given by $\mathbf{E}_p(\mathbf{x}, t) = \mathbf{E}_p \cos(\omega_p t - \mathbf{k}\cdot\mathbf{x})$, it interacts with the atoms in the electric dipole approximation:

$$\hat{H}_{\text{dip}} = -\hat{\mathbf{d}} \cdot \mathbf{E}_p \cos(\mathbf{k}\cdot\mathbf{x} - \omega_p t) \tag{2.15}$$

so that when we move into the probe-laser rotating frame and discard counter-rotating terms, the atoms evolve under a total Hamiltonian:

$$\begin{aligned}\hat{H}^{(R)} = \Bigg[&\sum_{\mathbf{n}} \Big(E_{\mathbf{n}} \left|g\right\rangle\left\langle g\right| \otimes \left|\mathbf{n}\right\rangle\left\langle \mathbf{n}\right| + (E_{\mathbf{n}} - \hbar\Delta) \left|e\right\rangle\left\langle e\right| \otimes \left|\mathbf{n}\right\rangle\left\langle \mathbf{n}\right| \Big) \\ &+ \sum_{\mathbf{n},\mathbf{m}} \left(\frac{\hbar\Omega_{\mathbf{nm}}}{2} \left|g\right\rangle\left\langle e\right| \otimes \left|\mathbf{n}\right\rangle\left\langle \mathbf{m}\right| + \frac{\hbar\Omega^*_{\mathbf{nm}}}{2} \left|e\right\rangle\left\langle g\right| \otimes \left|\mathbf{m}\right\rangle\left\langle \mathbf{n}\right| \right) \Bigg]\end{aligned} \tag{2.16}$$

with detuning $\Delta = \omega_p - \omega_0$ and with Rabi frequencies given by

$$\Omega_{\mathbf{nm}} = \Omega_0 \left\langle \mathbf{n}\right| e^{i\mathbf{k}\cdot\mathbf{x}} \left|\mathbf{m}\right\rangle \tag{2.17}$$

where $\Omega_0 = \left\langle g\right| \hat{\mathbf{d}} \cdot \mathbf{E}_p \left|e\right\rangle / \hbar$.

To evaluate the $\Omega_{\mathbf{nm}}$, this last expression can usefully be rewritten in terms of creation and annihilation operators:

$$\begin{aligned}\left\langle \mathbf{n}\right| e^{i\mathbf{k}\cdot\hat{\mathbf{x}}} \left|\mathbf{m}\right\rangle &= \prod_{i=1}^{3} \left\langle n_i\right| e^{ik_i\hat{x}_i} \left|m_i\right\rangle &(2.18)\\ &= \prod_{i=1}^{3} \left\langle n_i\right| e^{i\eta_i(\hat{a}_i + \hat{a}_i^\dagger)} \left|m_i\right\rangle &(2.19)\end{aligned}$$

where the η_i are the **Lamb-Dicke parameters** along each axis, given by:

$$\eta_i = k_i x_i = k_i \sqrt{\frac{\hbar}{2m\omega_i}} \tag{2.20}$$

Since the matrix elements factor out neatly into the three separate dimensions, we can just consider one dimension at a time, dropping the index i to avoid clutter. In this case, it is possible to show that [288][1]:

---

[1] The most obscure step in this derivation is to use the identity $e^{\hat{A}+\hat{B}} = e^{\hat{A}}e^{\hat{B}}e^{-[\hat{A},\hat{B}]/2}$. Otherwise we just expand out the exponential power series and act on $\left|n\right\rangle, \left|m\right\rangle$ with $\hat{a}, \hat{a}^\dagger$

$$\langle n| \, e^{i\eta(\hat{a}+\hat{a}^\dagger)} \, |m\rangle = \begin{cases} e^{-\frac{1}{2}\eta^2}\sqrt{\frac{m!}{n!}}\eta^{n-m}L_m^{n-m}[\eta^2] & \text{for red sidebands} \quad m < n \\ e^{-\frac{1}{2}\eta^2}L_n^0[\eta^2] & \text{for the carrier} \quad m = n \\ e^{-\frac{1}{2}\eta^2}\sqrt{\frac{n!}{m!}}\eta^{m-n}L_n^{m-n}[\eta^2] & \text{for blue sidebands} \quad m > n \end{cases} \tag{2.21}$$

where the $L_\beta^\alpha[\eta^2]$ are generalised Laguerre polynomials:

$$L_\beta^\alpha[\eta^2] := \sum_{j=0}^{\beta}(-1)^j \binom{\beta+\alpha}{\beta-j}\frac{\eta^{2j}}{j!} \tag{2.22}$$

Since the complexity of the Laguerre polynomials somewhat obscures what's going on, it is illustrative to consider the absorption signal from a weak probe beam. We keep to a 1D treatment, and assume that the atoms start in a thermal Boltzmann distribution:

$$P(n) = \mathcal{N}e^{-\frac{n_z\hbar\omega_z}{k_BT}} \tag{2.23}$$

$$\mathcal{N} = 1 - e^{-\frac{\hbar\omega_z}{k_BT}} \tag{2.24}$$

so that the absorption signal follows:

$$\begin{aligned} R_{sc}(\omega) \propto & \sum_{n<m} \mathcal{N}e^{-\frac{n\hbar\omega_z}{k_BT}} \frac{e^{-\eta^2}\frac{n!}{m!}\left(\eta^{m-n}L_n^{m-n}[\eta^2]\right)^2}{1+\frac{4^2}{\Gamma^2}(\omega-(\omega_0+(m-n)\omega_z))^2} \\ & + \sum_{n=m} \mathcal{N}e^{-\frac{n\hbar\omega_z}{k_BT}} \frac{e^{-\eta^2}\left(L_n^0[\eta^2]\right)^2}{1+\frac{4^2}{\Gamma^2}(\omega-\omega_0)^2} \\ & + \sum_{n>m} \mathcal{N}e^{-\frac{n\hbar\omega_z}{k_BT}} \frac{e^{-\eta^2}\frac{m!}{n!}\left(\eta^{n-m}L_m^{n-m}[\eta^2]\right)^2}{1+\frac{4^2}{\Gamma^2}(\omega-(\omega_0+(m-n)\omega_z))^2} \end{aligned} \tag{2.25}$$

This equation is the basis for the numerical calculations in figure 2.21, where we depict the absorption spectrum for a large range of trap strengths ω_z. As the confinement increases in strength, we can observe a series of regimes:

- $\omega_z < \Gamma$: A regime of weak confinement, where the fluoresence resembles a free-space Gaussian spectrum with Doppler-broadened width $\delta\omega = kv_T$
- $\omega_z > \Gamma$: The **resolved-sideband regime**, where quantized motional transitions become distinguishable under the Gaussian envelope

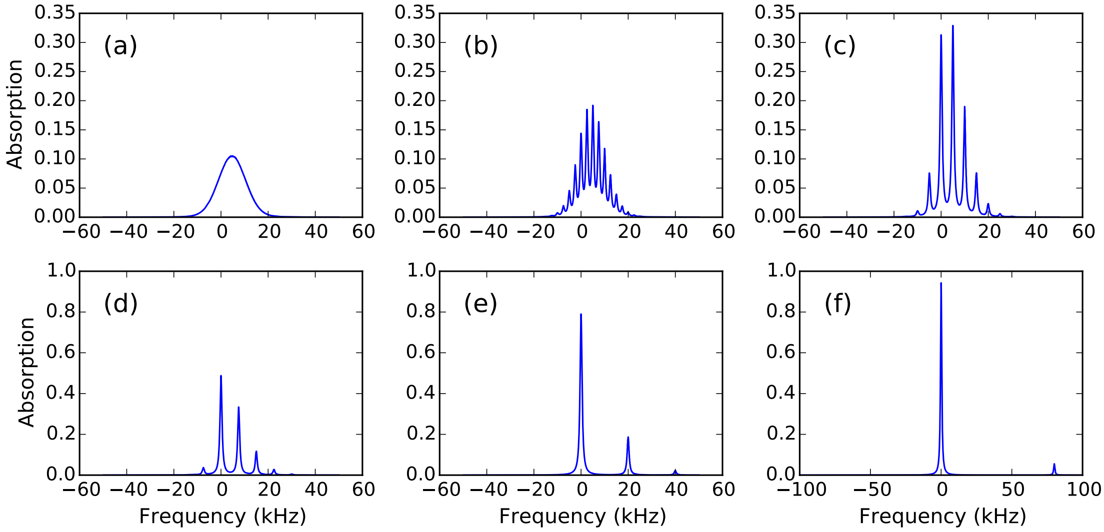

Figure 2.21: The modelled absorption spectrum from 1 µK atoms at various strengths of confinement in a 1D harmonic potential: (a) 500 Hz, (b) 2.5 kHz, (c) 5 kHz, (d) 7.5 kHz, (e) 20 kHz, and (f) 80 kHz. For illustrative purposes we introduce a decay rate of $\Gamma = 2\pi \times 1\,\text{kHz}$. The y axis indicates the weak-probe absorption relative to the on-resonance absorption of a fixed-position atom.

- $\omega_z > \omega_{rec}$: The **Lamb-Dicke regime** ($\eta < 1$), where the recoil shift no longer dominates and the carrier becomes the strongest transition
- $\hbar\omega_z > k_B T$: The **ground-state cooled regime**, where the sidebands are further suppressed, especially on the red-detuned side

Among these different regimes, the numerical modelling reveals that the best situation for clock spectroscopy is in the limit of strong confinement, where the motional sidebands are most strongly suppressed and where the carrier strength asymptotically reaches that of the 'ideal' motionless atom used in the definition of the SI second (see section 1.1.1).

To gain more insight into this regime, and returning now to a full 3D treatment, we can Taylor expand $e^{i\eta_i(\hat{a}_i+\hat{a}_i^\dagger)}$ in the limit of small η_i. Keeping only the 1st-order terms in η_i, we can show that the carrier Rabi frequency simplifies to:

$$\Omega_{\mathbf{nn}}/\Omega_0 = 1 - \sum_{i=1}^{3} \eta_i^2 \left(n_i + \frac{1}{2} \right) + \mathcal{O}(\eta_i^4 n_i^2) \tag{2.26}$$

thus confirming carrier enhancement $\Omega_{\mathbf{nn}} \to \Omega_0$ as $\eta_i^2 \to 0$.

The nearest red and blue sidebands are the only other features that remain to first order in η_i, with:

$$\begin{aligned} \Omega_{\mathbf{n},\mathbf{n}+\mathbf{i}}/\Omega_0 &= \eta_i\sqrt{n_i+1} \;+\; \mathcal{O}(\eta_i^3 n_i^{3/2}) \\ \Omega_{\mathbf{n},\mathbf{n}-\mathbf{i}}/\Omega_0 &= \eta_i\sqrt{n_i} \qquad\; +\; \mathcal{O}(\eta_i^3 n_i^{3/2}) \end{aligned} \tag{2.27}$$

thus confirming sideband suppression $\Omega_{\mathbf{n}\neq\mathbf{m}} \to 0$ as $\eta_i\sqrt{n_i} << 1$.

As well as explaining the dominance of the carrier transition in the Lamb-Dicke regime, equations 2.27 also help to quantify the asymmetry of the fluorescence spectrum in the ground-state cooled regime. As we reduce the atom temperature to $k_B T \ll \hbar\omega_i$, the average excitation n_i diminishes exponentially as:

$$\bar{n}_i = \frac{\sum_{n_i} n_i e^{-\frac{n_i\hbar\omega_i}{k_B T}}}{\sum_{n_i} e^{-\frac{n_i\hbar\omega_i}{k_B T}}} = \frac{1}{e^{\frac{\hbar\omega_i}{k_B T}} - 1} \longrightarrow 0 \tag{2.28}$$

and thus the red sideband is suppressed as $\sqrt{n_i} \to 0$, but the blue sideband still remains as $\sqrt{n_i+1} \to 1$. This bias towards absorption of blue-detuned light is the trapped-atom analogue of the free-space recoil effect.

To conclude, equations 2.26 and 2.27 show that we can suppress the Doppler and recoil effects by tightly confining the atoms such that $\eta_i\sqrt{n_i} << 1$ in all directions. This requirement places one important constraint on the geometry of optical lattice clocks: since the recoil frequency of the 698 nm Sr clock transition is 4.7 kHz – much higher than radial trap frequencies, but comfortably lower than a typical axial trap frequency – we can only achieve small η_i when the atoms are probed along the lattice axis, i.e. when $\mathbf{k}_p$ is mostly pointing along $\mathbf{k}_L$. However, as long as we stick to this specific geometry, we can resolve the narrow Doppler- and recoil-free central features like those in figure 2.22 – an ideal basis for making a stable, accurate strontium optical lattice clock.

2.3.3.3 The magic wavelength

The optical lattice is a excellent way to confine the atoms, but it also threatens to introduce a significant systematic clock error: the AC Stark shift to the clock transition. Indeed, at first sight this shift seems prohibitively large: we aim for clock accuracy at the few mHz level, but even to make a shallow trap with a depth of 1 µK we must induce a ground-state Stark shift of more than 20 kHz.

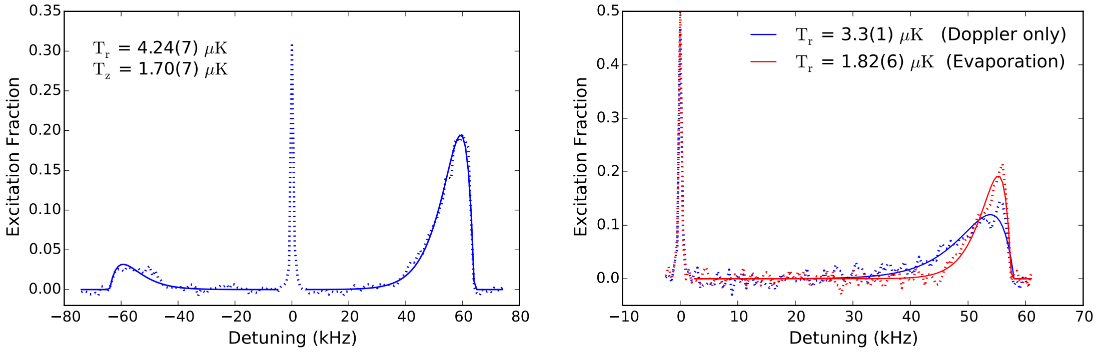

Figure 2.22: *Left:* Scans over the narrow carrier transition and the broader axial lattice sidebands (dotted) with theoretical fits (solid) based on eq 2.18 from [282]. The radial temperature is inferred from the sideband shape, which is broadened for hotter atoms due to trap anharmonicity; the axial temperature is calculated from the ratio of blue and red sideband amplitudes. *Right:* Using an evaporation stage (ramp down to 2 µK lattice depth in 27 ms, hold for 27 ms, then ramp back up in 27 ms) we observe a narrower blue-detuned sideband, indicating that only the colder atoms remain in the trap.

Thankfully there is a viable solution at hand. The *overall* frequency shift to the clock transition from the lattice is given by:

$$\delta\nu_{\mathrm{St}} = -\frac{1}{4h}\left(\alpha_e(\nu) - \alpha_g(\nu)\right)|\mathbf{E}(\mathbf{x})|^2 \tag{2.29}$$

where $\alpha_g(\nu)$ and $\alpha_e(\nu)$ are the polarisabilities of the ground and excited states respectively[1]. The key insight, originally proposed as the basis for optical clocks in [128], is that by choosing a particular "magic" wavelength we can engineer a differential polarisability $\Delta\alpha(\nu) = \alpha_e(\nu) - \alpha_g(\nu) = 0$. At this wavelength, the net shift to the clock transition frequency is (nearly) zero regardless of the intensity of the lattice light, greatly relaxing the constraints on accurate clock spectroscopy.

In Sr we observe viable magic wavelengths at 390 nm, 497 nm, and 813 nm, as highlighted in the polarisability plot in figure 2.23. Historically, and in this thesis, the crossing at 813 nm has been used for lattice clocks because of: (1) the availability of high-power laser sources, (2) the low scattering rate compared with 497 nm, and (3) the relative convenience of operating a red-detuned dipole trap in contrast with the blue-detuned 390 nm trap. However, the low power requirements at 497 nm or the close lattice spacing at 390 nm might give us cause to swap wavelengths in the future.

---

[1] In this formula we use the electric field amplitude $\mathbf{E}(\mathbf{x})$, which is a factor $\sqrt{2}$ larger than the time-averaged r.m.s. electric field $\langle\mathbf{E}(\mathbf{x})\rangle_{\mathrm{rms}}$.

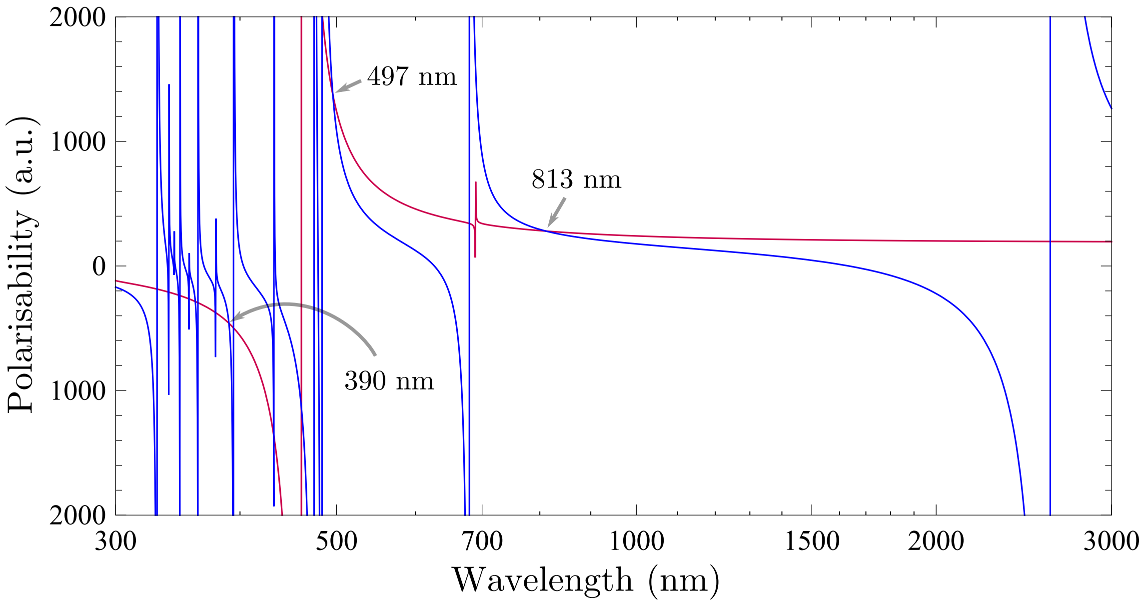

Figure 2.23: Theoretical polarisabilities of the ground (red) and excited (blue) clock states over a range of wavelengths, calculated by summation of contributions from second-order perturbation theory using a program developed by Ian Hill [106]. The arrows highlight three favourable magic wavelengths, of which two have so far been experimentally verified [193, 260].

2.3.3.4 The lattice laser system

When I joined the project, there was already an all-semiconductor laser system in place to supply the light for the lattice trap, though the trap itself hadn't yet been set up. An AR-coated 820 nm Toptica laser diode was integrated into an ECDL as described in [106], and its 20 mW output was used to seed a Toptica 'BoosTA' tapered amplifier nominally outputting 500 mW at 813 nm. Unfortunately, this laser system turned out not to be a great solution for a lattice clock, mostly because of the TA chip:

1. The TA output power and beam quality was too low, such that after the isolator, the AOM and the fibre we only had around 200 mW left to make the lattice trap at the atoms.

2. The pedestal of amplified spontaneous emission (ASE) from the TA covered a very broad ∼ 30 nm spectrum, causing significant AC Stark shifts from non-magic frequency components. This could have been partially solved by installing an interference filter [284], but only at the expense of further power losses.

3. The TA chip itself was too prone to failure – it broke after only a couple of months of use, and, quite frustratingly, almost immediately after we got the first atoms into the lattice trap for the first time.

Because of these limitations, and because we were not the only group to have reliability problems with this chip, we decided not to buy a like-for-like replacement. Instead, we purchased a SolsTiS laser from M Squared Lasers. This Ti:Sapphire laser has plentiful 6 W output power, excellent beam quality, very low ASE, a narrow few-kHz free-running linewidth, and shows no sign of deterioration after 2 years of use. The only downsides are a slightly increased mechanical sensitivity (tapping the table can often unlock the laser) and a price tag that some taxpayers may find offensive. In future systems aiming at lower size, weight, power, and cost (SWaP-C), the best way forward will probably be to send the output of an 813 nm ECDL or DFB straight to a power enhancement cavity around the atoms, perhaps with extra interference filtering if necessary.

With the new commercial Ti:sapphire laser, the rest of the lattice setup becomes pretty simple. At the laser's output, around 0.1% of the light is picked off using a backside-polished mirror and is sent to the transfer cavity through a double-passed AOM and an EOM to stabilise the laser's frequency to within roughly 10 kHz (see section 2.2.1). The rest of the light (up to 6 W) is sent through an optical isolator and split two ways for use at two different positions around the science chamber – one path is for the main vertical lattice trap, and the other is either for a horizontal lattice or for 3-photon spectroscopy (see section 4.2.3). Each of these two paths is sent through a single-passed AOM for intensity stabilisation, and is then coupled into separate high-power-connectorised polarisation-maintaining fibres from OZ optics (part numbers PMJ-3AHPC,3A-633-4/125-3AS-5-1 and PMJ-3HPC,3HPC-810-5/125-3AS-5-1-AR2-CF2). We have never observed damage of either fibre up to 1 W fibre-coupled power, and we have found no evidence of non-linear effects such as Brillouin scattering when looking at atomic heating rates, clock-transition frequency shifts, or at optical beat signals between free-space and fibre-coupled beams.

The optical setup for the main vertical lattice trap is depicted in figure 2.24. We focus down to a 65 µm waist[1] at the centre of the science chamber using a single Thorlabs AL1225-B aspheric lens in an SM1ZM zoom housing screwed into the same LMR1/M mount as holds the fibre-launching adapter. After passing through the science chamber, the lattice beam is recollimated on a cage-mounted 500 mm

---

[1] In an earlier setup the waist used to be 45 µm, so don't be confused by any pre-2016 publications.

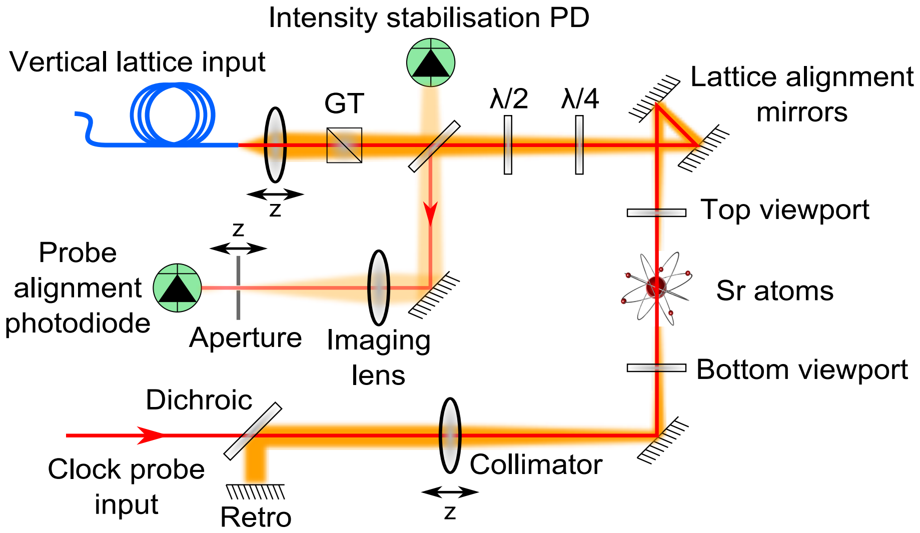

Figure 2.24: The complete optical system for the vertical lattice trap. The 813 nm lattice beam is actually a dim red to the eye, but orange is used here to distinguish from the red 698 nm clock probe beam. "GT" denotes a Glan-Taylor polariser.

achromatic-doublet lens, then combined with the clock probe beam on a short-pass Semrock FF750-SDi02 dichroic, and finally retro-reflected on a dielectric mirror in a FineAdjust Mercury-series kinematic mount from Radiant Dyes. We make sure to combine the clock probe beam physically very near the lattice retro-reflector so that the sampler for clock path-length stabilisation (not shown in the figure) can be as close as possible to the retro-reflector; this way, the total uncompensated path can be kept to less than 10 cm so as to mitigate probe phase noise (see section 3.1.1).

On the return path for the lattice beam, the same beam sampler used in the first pass for intensity stabilisation now diverts some light into a rudimentary single-lens (Thorlabs AC508-150-B-ML achromat) imaging system to allow straightforward tweaking of the clock probe beam alignment. This tweaking is done by the following procedure:

1. (Only necessary once) Walk the z-translations of the fibre-launch and lattice-collimating lenses to maximise the axial sideband frequency for a fixed input lattice power, taking care to re-optimise retro-reflector alignment at every step. This ensures that the in-going and retro-reflected lattice-beam waists are both focused at the atoms.
2. Walk the pre-imaging-lens mirror (the one right in the middle of figure 2.24) and the aperture z-translation to maximise the signal from the lattice beam on

the probe alignment photodiode. This ensures that the aperture is centred at the lattice waist in x, y and z, and therefore images the lattice-trapped atoms.

3. Finally, align the clock beam on one mirror[1] to maximise the signal on the probe alignment photodiode; the beam is then passing through the aperture with the lattice, and thus also through the lattice-trapped atoms.

We have found this alignment procedure to be very effective, as verified through observations of the probe AC Stark shift in ^{88}Sr. As a result, the vertical lattice system quite easy to maintain; after a one-off complete optimisation, only two things have needed regular tweaking: (1) the lattice retro-reflecting mirror, by maximising either the back-fibre-coupling signal or the axial sideband frequency, and, less often, (2) the clock probe alignment, by the procedure outlined above.

With this apparatus in place, along with the many cooling lasers described throughout this chapter, we can reliably fill the lattice trap with more than 10^6 atoms of ^{88}Sr or 10^4 atoms of spin-polarised ^{87}Sr in around 300-500 ms. As we elaborate in the next chapter, the next step is to probe the clock transition and to read out the atoms' response with the ultimate goal of stabilising the optical clock.

---

[1] Rough probe beam alignment should first be done on two mirrors, either by eye or using coupling efficiency into the lattice-launch fibre, to avoid any nasty effects of radial recoil [43].

Chapter 3

Clock stability

One of the most impressive features of optical atomic clocks is their exceptional stability. In the most stable clock published, we'd expect the frequency to tick at the same rate to 18 decimal places from one 7-hour period to the next [6], allowing us to detect fractional frequency changes at the minuscule level of 0.000000000000000001. Since there are no signs of a fundamental source of drift, optical clocks can be used to measure with even more digits of precision by extending the averaging time further. Even so, making optical clocks more stable is still an important research priority; if we are ever going to exploit their 10^{-18} accuracy in practical applications, it will help to reduce the required measurement times to a few minutes rather than several hours.

In this chapter, we explore in detail the layers of frequency control needed to stabilise a Sr optical lattice clock. As we sketch in figure 3.1, different control loops take over at different Fourier frequencies of noise; the high-bandwidth 'prestabilisation' control loops will dominate the faster clock frequency fluctuations, but the ultimate long-term frequency stability (and, as it happens, the ultimate clock accuracy) arises from the slowest lock to the atomic reference.

Throughout this chapter we make extensive use of the Allan deviation $\sigma_y(\tau)$ as a way of characterising clock instability, defined by [226]:

$$\sigma_y^2(\tau) \equiv \frac{1}{2}\left\langle \left(\frac{1}{\tau}\int_{t+\tau}^{t+2\tau} y(t')dt' - \frac{1}{\tau}\int_{t}^{t+\tau} y(t')dt' \right)^2 \right\rangle \tag{3.1}$$

where $y(t) := (\nu(t) - \nu_0)/\nu_0$ is the fractional frequency error, τ is the 'averaging time' over which stability is measured, and the overall brackets $< \cdot >$ describe an average over all start times t. The Allan deviation has a quite straightforward interpretation: if we split the frequency data $y(t)$ into time bins of length τ, the Allan deviation is a measure of the expected difference in average $y(t)$ between neighbouring bins. Note

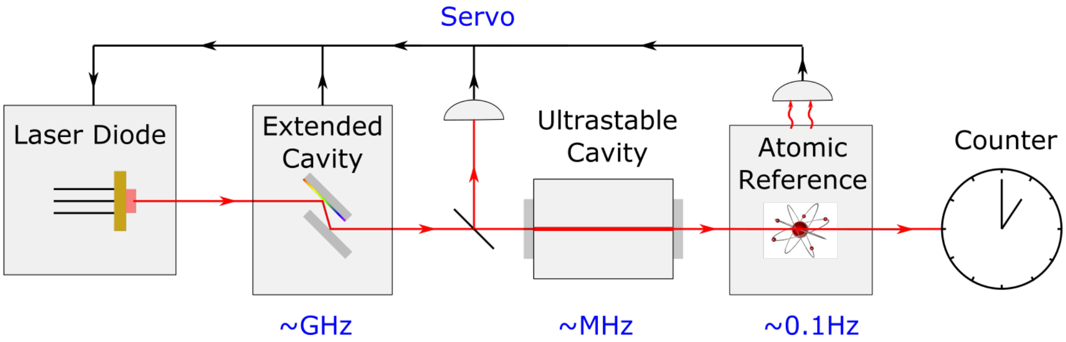

Figure 3.1: A diode laser is made into an optical clock using multiple feedback loops of different bandwidths (shown in blue text). In the jargon of the clock-making community, we group together all of the control loops except for the final one – the lock to the atomic reference – as the 'local oscillator'.

that for white frequency noise, the neighbouring bins are uncorrelated, so the Allan deviation is equal to the true standard deviation – see [226] for further discussion.

We can also describe clock instability in the frequency domain, using one-sided power spectral density $L_y(f)$ or two-sided power spectral density $S_y(f)$:

$$S_y(f) := \lim_{T\to\infty} \frac{1}{T}\left|\int_{-T/2}^{T/2} y(t)e^{2\pi i f t}\,dt\right|^2 \tag{3.2}$$

$$L_y(f) := \begin{cases} 2S_y(f) & \text{for } f > 0 \\ S_y(f) & \text{for } f = 0 \\ 0 & \text{for } f < 0 \end{cases} \tag{3.3}$$

The time-domain representations of various common types of noise spectrum are given in table 3.1 and depicted in figure 3.2 (conversions are taken from [226]).

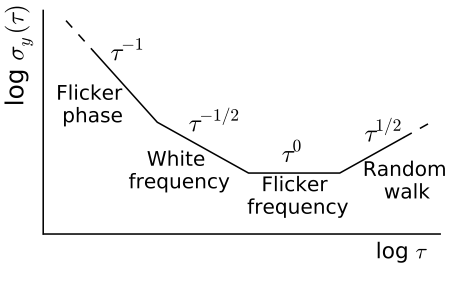

Noise type	$L_y(f)$	$\sigma_y(\tau)$
white & flicker phase	f^2 & f	$1/\tau$
White frequency	h_0	$\sqrt{h_0/2\tau}$
Flicker frequency	h_1/f	$\sqrt{2h_1 \ln 2}$
Random walk	h_2/f^2	$\sqrt{\frac{2}{3}\pi^2 h_2 \tau}$

Figure 3.2 & Table 3.1: White and flicker phase noise both lead to a scaling $1/\tau$ for $\sigma_y(\tau)$, but the constant of proportionality depends on measurement bandwidth.

3.1 The local oscillator

A robust, narrow-linewidth local oscillator is a critical component of a clock, not only determining its short-term frequency noise but also feeding into its asymptotic long-term stability at averaging times $\gtrsim 10\,\mathrm{s}$. As we explore in detail in section 3.3, there are two important mechanisms by which the local oscillator affects long-term clock stability: first (dominant in ion clocks), the atomic interrogation time is often limited by the coherence time of the interrogating laser, imposing a high quantum projection noise instability if the laser decoheres too quickly [121, 209]. Second (dominant in optical lattice clocks), short-term laser frequency fluctuations are usually sampled unevenly by the atoms, aliasing into long-term atom lock instability through the Dick effect [69]. Since the long-term stability limits how fast we can make accurate frequency measurements, a better local oscillator makes for a much more useful clock.

In our Sr clock, we follow the established method for building high-performance optical local oscillators, stabilising a single-mode laser to an ultrastable reference cavity using the Pound-Drever-Hall (PDH) technique [73]. This approach has yielded the narrowest-linewidth lasers yet built [99, 135], though further advances in stability have become increasingly challenging (and expensive) as the technology has matured. In this section we describe our 698 nm clock laser, highlighting the main sources of instability which limit its performance as a local oscillator.

3.1.1 Overview of the clock laser system

A simplified schematic of the clock laser system is shown in figure 3.3, and more details are depicted in figure 3.4. The centrepiece is the reference cavity housed inside the vacuum chamber described in [46], which is then mounted on an active vibration-isolation (AVI) platform and sandwiched from above and below by a pair of optical breadboards densely populated with optics. The whole system sits on the corner of a room-sized, passively damped optical table, and is enclosed in a $60 \times 63 \times 75$cm transparent box to shield from air flows and temperature fluctuations.

The source of 698 nm clock light is a non-AR-coated Opnext HL7001MG diode. Despite a nominal centre wavelength of 705 nm, we've conveniently found that the output of these diodes tends to be within a couple of nm of the Sr clock frequency at room temperature. However, the free-running InGaAsP diode has a broad linewidth on the order of 100 MHz due to its short, low-finesse cavity, making it unsuitable for coherent atomic interrogation. To tame the linewidth to below 100 kHz, we place the diode into an extended lasing cavity.

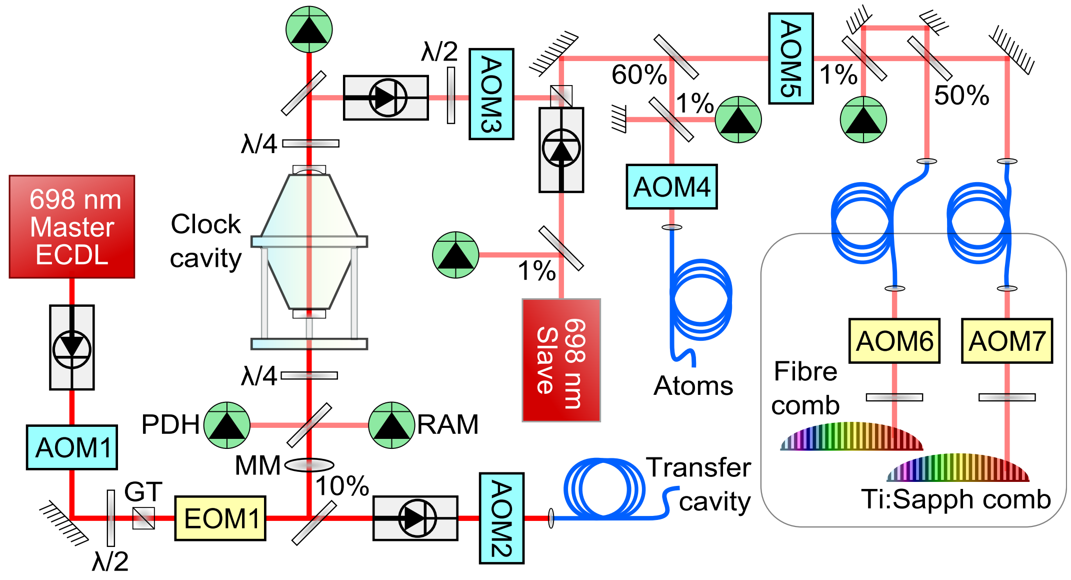

Figure 3.3: Lots of AOMs are needed for clock frequency and intensity control. AOM1: Cavity intensity stabilisation. AOM2: Frequency bridge to transfer cavity. AOM3: Atom-lock frequency corrections and drift compensation. AOM4: Path-length stabilization (PLS) to atoms and universal synthesizer feedforward. AOM5: Frequency bridge to combs. AOM6 & AOM7: PLS to combs.

Our home-built Littrow-configuration extended-cavity diode laser (ECDL), shown in figure 3.5, has been described in detail in the two previous doctoral theses from the project [46, 106]. Designed by Ian Hill, much attention was paid to make the ECDLs robust and easy to use: (1) a monolithic hermetically-sealed configuration reduces alignment drift and supplies good immunity to background vibrations and pressure changes; (2) a well-chosen pivot point for the diffraction grating allows us to scan the cavity modes in equal measure to the peak diffraction wavelength, increasing mode-hop-free tuning into the several GHz range [197]; (3) a mirror integrated into the grating mount avoids any significant changes in the alignment of the output beam while tuning the grating angle [103]; and (4) good temperature control of the whole lasing cavity enables stable continuous operation for several days without mode-hops. After setting up the diode in its extended cavity, we achieve a much reduced short-term linewidth of around 70 kHz and fairly slow drift of tens of MHz per hour [46]. Of course, this still isn't good enough to perform Hz-level clock spectroscopy, so we have to apply another layer of frequency stabilisation: the PDH lock to the reference cavity.

Our vertical 7.75 cm reference cavity is depicted in figure 3.5, with both the spacer and the mirror substrates constructed from Corning ultra-low expansion (ULE) glass.

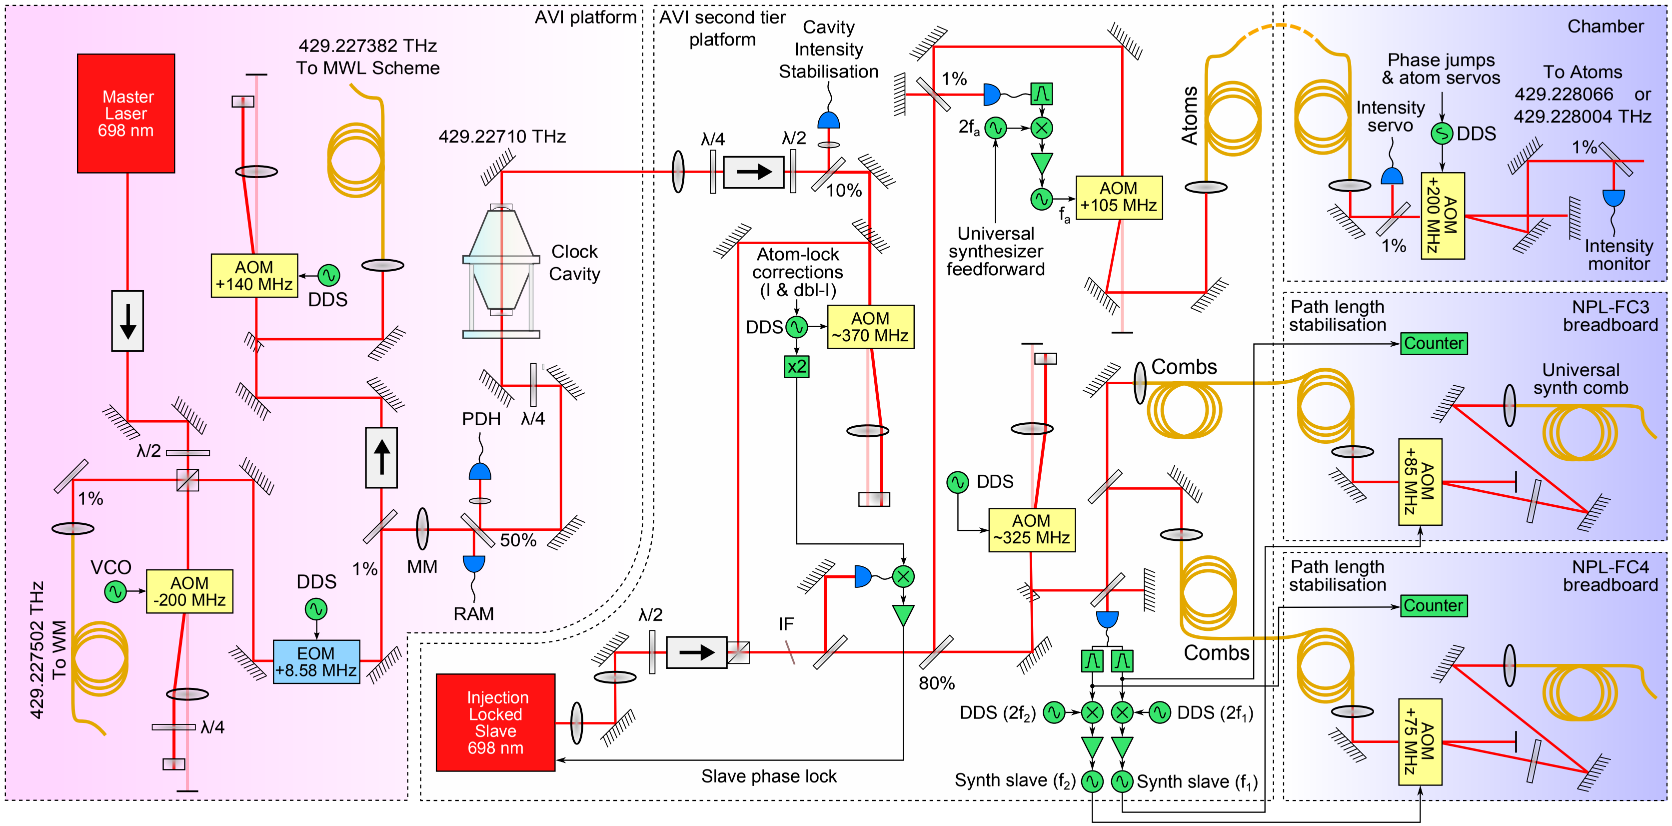

Figure 3.4: A more detailed schematic of our 698 nm laser system.

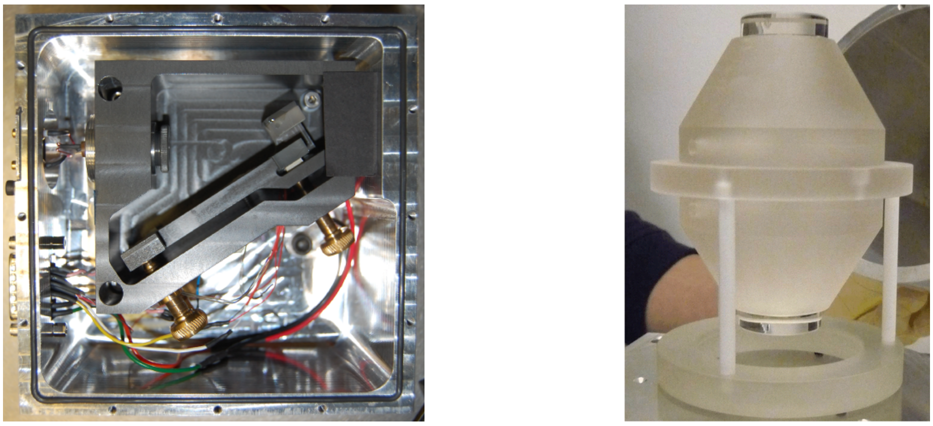

Figure 3.5: Pictures of our home-built 698 nm ECDL (left) and the 'American football' reference cavity (right).

The mirrors are coated with a SiO_2/Ta_2O_2 multilayer Bragg stack, giving us a $420,000$ cavity finesse with 5 ppm mirror transmission and 2.5 ppm round-trip loss at 698 nm (measured in [46]). We use a home-built loop filter (see section 2.2.3) in the PDH lock, applying fast, DC-limited, singly-integrated feedback to the diode current for narrowing the short-term laser linewidth, and slow doubly-integrated feedback to the piezo to compensate the long-term drift in ECDL length. After optimizing the loop gains, we observe a cavity transmission which is flat well below the 1% level, heuristically indicating that the laser linewidth must be much less than the cavity linewidth of 4.5 kHz. We show that PDH lock noise is sub-Hz in the next section.

In order to filter out the PDH servo sidebands (making it easier to count the beat signals against the Ti:Sapphire frequency comb), we use the cavity transmission to injection-lock another uncoated HL7001MG slave diode. With quite substantial power amplification needed in typical operating conditions – from around 1 µW input to 15 mW output – the slave can be quite fiddly to work with. To inject the slave properly (i.e. to stop it from pulling away to another frequency) it is critical to get its current and temperature exactly right, so that one of the longitudinal modes of the slave diode rests at the injection frequency. Even when a mode is lined up, we find that the slave will occasionally glitch if operated at $\gtrsim 20$ mW. (We detect this as a spike in very broadband noise on our path-length-stabilisation photodiodes). At typical operating conditions, and with optimized alignment, we observe an injection window of around 150 µA, just about enough to stay properly injected over a few days. We observe a highly asymmetric behaviour of the slave with more glitches progressively emerging towards the lower-current end of the injection window, in good agreement with the expected effect of changing diode refractive index upon injection [105] –

because of this, we usually park the slave current at the higher end of the window.

It's worth mentioning our attempt to phase-lock the slave using the configuration depicted in figure 3.4 (see the components around the 370 MHz AOM towards the middle of the figure). The idea is to inject with the two frequencies generated by the double-passed AOM: a more intense frequency component picks up a 740 MHz shift by getting diffracted twice by the AOM, while a less intense component passes twice through the AOM without getting diffracted. Both co-propagating beams are then aligned into the slave diode, and we centre the slave current to be injected well by the more intense frequency component. To implement the lock, we sample the light coming out of the slave on a fast photodetector, and then we lock the phase of the resulting 740 MHz beat to a doubled version of the 370 MHz AOM drive frequency, actuating on the slave diode current. The idea is that if the slave mode drifts upwards in frequency it will develop a phase lead compared with the injected light, and if it drifts downwards it will develop a phase lag compared with the injected light. These optical phase changes will feed onto the phase of the detected RF beat, and give us an error signal for servoing the slave back onto resonance. It worked! However, the long-term stability was undermined by drift in the relative path length of each frequency component: over the course of a day or so the paths were clearly fluctuating by a significant fraction of an optical wavelength, enough to render the phase-lock useless. In the end, the stability of the free-running slave turned out to be sufficient, and so the phase lock apparatus has remained vestigial. If we ever require a phase lock in the future, direct frequency modulation of the AOM drive may prove to be more robust against path-length drift.

Even after building the rather stable master and slave lasers described above, there's a lot more work needed to distribute the light where it's needed (see all the optics in figure 3.4). Within the maze of AOMs we have three copies of the type of path-length stabilization setup first introduced in [170]; these setups detect and compensate for any changes in optical path sampled by the clock beam (primarily from vibrations and temperature fluctuations which modulate the refractive index of optical fibres), preventing excess phase noise that would otherwise undermine atomic interrogation or combs frequency measurements. The various other AOMs in the clock apparatus perform fairly standard tasks of intensity stabilisation, frequency correction, and spectroscopic pulse shaping.

A final step to improve clock laser stability is the frequency feed forward from the so-called universal synthesizer, tying the clock light used for atomic interrogation to an even more stable laser at 1064 nm sitting in another lab down the corridor. Since

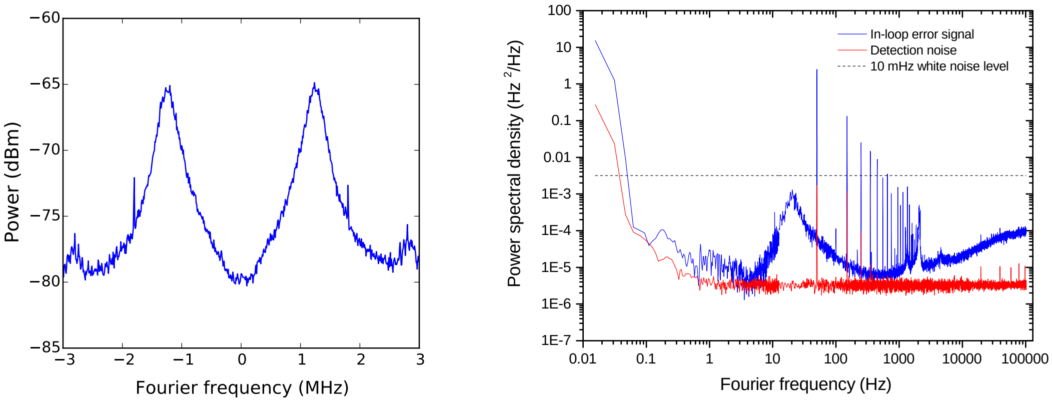

Figure 3.6: A spectrum analyser trace around the 8.6 MHz beat frequency from the clock PDH photodetector showing servo sidebands at ∼ 1.3 MHz (left). Fast Fourier transforms of the error signal, normalised to power spectral density of laser frequency fluctuations (right). (Note that the normalization is only valid within the cavity linewidth of 4.5 kHz, outside which the error signal acts as a discriminant for phase rather than frequency).

this was only implemented as an add-on to the clock system in late 2015, we leave a detailed discussion of the universal synthesizer to section 3.1.3.

3.1.2 Limits to local oscillator stability

Good discussions of the sources of instability in cavity-stabilised lasers can be found in various other doctoral theses [164, 175, 282]. In this section, we run through some of the effects that influence our local oscillator performance.

3.1.2.1 PDH lock noise

As discussed in section 2.2.1.1 in the context of the transfer cavity scheme, the dominant optical source of PDH lock instability is often residual amplitude modulation (RAM), especially at low Fourier frequencies. Fortunately, our high-finesse clock cavity has a relatively small linewidth of 4.5 kHz, so we only need to resolve the central resonance to within around 10^{-4} – a lot easier than the 10^{-6} level achieved with some effort in [303]. In practice we have found that after careful tuning of the pre-EOM Glan-Taylor polarizer, RAM-induced instability in the 698 nm system stays at a negligible level.

After RAM, other imperfections in the PDH control loop can become significant. In practice, the main consideration here is to achieve sufficient loop gain to suppress the free-running laser frequency noise to the required level. We use two methods to diagnose whether the lock is performing well enough (see figure 3.6). In the first method we take fast Fourier transforms of the PDH error signal, both when the laser is in lock and when it's out of lock. To convert the error signal noise into laser frequency fluctuations we use a discriminant $dV/d\nu$ calculated from the measured peak-to-peak mixer output and the known linewidth of the cavity. In the locked trace we observe a small noise peak at around 20 Hz (possibly a servo hump from the slow PZT feedback), but the overall noise is negligible at well below $3.1 \times 10^{-2}\,\mathrm{Hz}^2/\mathrm{Hz}$. White noise at this level would correspond to a laser linewidth of $\ll 10\,\mathrm{mHz}$, much less than our expected frequency instability from cavity length fluctuations. The unlocked trace reveals the detection noise in the setup, again much lower than we need to worry about.

The second PDH-loop diagnostic is the power spectrum of the PDH photodetector signal around the EOM modulation frequency of 8.6 MHz. This signal is especially important when optimizing the PDH servo loop since it show us the noise at Fourier frequencies around the unity-gain bandwidth of the lock, where free-running laser fluctuations are barely suppressed (or can often be enhanced) by the loop. An important aim in the optimization procedure is to achieve a high servo bandwidth, since we observe that the laser will quickly come unlocked unless the servo sidebands are pushed out well beyond 500 kHz from the carrier. We also use the spectrum analyser signal to avoid servo oscillations, detected through the emergence of "spikier" peaks on the servo sidebands − if oscillations are allowed to persist we find that the cavity transmission is less stable, presumably because the amount of optical power shunted into the servo sidebands becomes vulnerable to varying loop gain.

3.1.2.2 Minor effects: pressure, mirror heating, material creep

In an ideal PDH lock, the laser frequency perfectly tracks changes in the cavity resonance. To a good approximation, fractional fluctuations of cavity length and laser frequency are then exactly the same, giving us the relation:

$$\frac{\Delta\nu}{\nu_0} = \frac{\Delta L}{L} \tag{3.4}$$

In this brief section we discuss a few causes of cavity length fluctuations which we label as 'minor', but which can lead to significant frequency instability if appropriate design precautions aren't taken.

First we consider the pressure of gases inside the cavity, which can change the effective cavity length through its refractive index. In the low-pressure limit with 698 nm light propagating through dry air at 25 °C, the pressure shift is

$$\frac{\Delta\nu}{\nu_0} = \frac{\Delta n}{n_{\text{vac}}} = 2.6 \times 10^{-7} P \tag{3.5}$$

where P is the pressure in mbar [133]. To keep this effect small we place the cavity in vacuum at a pressure of 2×10^{-8} mbar, leading to a 5×10^{-15} fractional frequency shift compared with perfect vacuum. We can therefore tolerate 10% pressure fluctuations before seeing significant changes in frequency – in practice, we see no evidence of a pressure spikes at this level.

The next consideration is the high laser power circulating in the cavity. Our cavity has $420,000$ finesse, implying a factor of $134,000$ enhancement between the amount of light coupled in and the steady-state power actually circulating inside. Since the cavity transmission needs to be high enough reliably to inject a diode slave laser, we typically couple between 5 and 10 µW in. This may not sound like much, but it gets enhanced to around 1 W bouncing between the mirrors. In theory, even the resulting radiation pressure could be significant at the 10^{-16} level [282]. However, the larger light-induced shift is from localized mirror heating – sensitivities ranging from 4 to 40×10^{-15}/µW of input power have been observed in similar systems [23]. In our case, we implement a long-term active intensity stabilisation loop to prevent drifts in circulating power – we find that this significantly improves the linearity of long-term cavity drift. At short timescales, the passive power stability is sufficient not to ingress at the 10^{-15} level.

Finally, we have to take into account the linear drift in cavity length as the spacer material ages. This effect is especially large for amorphous spacer materials like ULE, in our case resulting in a drift rate of 95 mHz s^{-1} (see figure 3.7). Fortunately the drift rate is quite steady, so we can compensate for it by feeding a linear ramp onto the frequency of an AOM later in the beam path. After compensation is applied, the drift no longer threatens to affect the local oscillator coherence time or the noise spectrum, though as a precaution against possible offsets we do have to apply double integration in the control loop when locking to the Sr clock transition (see section 3.2.1).

3.1.2.3 Vibrations and temperature fluctuations

The frequency stability of cavity-locked lasers is quite remarkable considering the physical dimensions involved: even our relatively wobbly cavity is stable in fractional length to roughly 1×10^{-15} at 1 s, meaning that the mirror separation of 7.75 cm remains fixed to within 0.1 fm over that averaging time, i.e. 1/10 of the proton radius! Unsurprisingly, the effects of mechanical vibrations and thermal expansion can become significant at this level, presenting some of the most important considerations in cavity design.

As originally modelled [52] and tested [8, 166, 201] at JILA, the rationale for mounting the cavity in a vertical orientation is to provide good immunity of cavity length against vibrations: roughly speaking, any upward force transmitted onto the mounting points at the midplane of the cavity will symmetrically compress the top half and stretch the bottom half, leaving the net distance between the mirrors unaffected.

Our ULE vertical cavity has exactly the same vertical design as one of those previously used at JILA, where they measure quite a low vibration sensitivity of $7 \times 10^{-10}/\mathrm{g}$ along the most vulnerable direction [166]. (Here we use units of fractional length change relative to acceleration under gravity $g = 9.81\,\mathrm{m\,s^{-2}}$). However, we probably get worse vibration performance from the same cavity since we haven't gone through any rigourous process of measurement or optimization. As a result, even after applying active and passive vibration isolation, we found out in late 2015 that there is a large noise spike on our clock laser at a Fourier frequency of around 2 Hz (see figure 3.10). We can confidently attribute this noise spike to mechanical vibrations by observing very large fluctuations at this frequency when we even so much as lightly prod our optical table; clearly the oscillations are from the optical table settling (or perhaps the AVI). We could potentially work towards eliminating this noise source, either by remounting the cavity to reduce sensitivity or by using some sort of vibration detection and active feedforward technique like the one in references [149, 265], but fortunately we now have the luxury of a universal synthesizer which can detect the frequency fluctuations directly and lets us compensate accordingly (see section 3.1.3).

As a brief digression, we note that vibration-insensitive alternatives to the vertical design are gaining favour in the latest generation of optical cavities. In particular, rigidly-mounted cubic [278] and spherical [150] cavities give the best reported performance at $2 \times 10^{-11}/\mathrm{g}$ along the worst axis, making them especially suitable for space-based or portable clocks. Meanwhile, horizontal cavities are being used as state-of-the-art laboratory references since they are much more practical to extend to lengths greater than 30 cm needed to reduce thermal noise (as we discuss in the next section). With careful mounting, the horizontal cavities can still demonstrate low $1.7 \times 10^{-10}/\mathrm{g}$ sensitivity even at an unwieldy length of 48 cm [99]. Nonetheless, our compact vertical design has found its niche in some cryogenic systems such as [135] where space within the cryostat is at a premium.

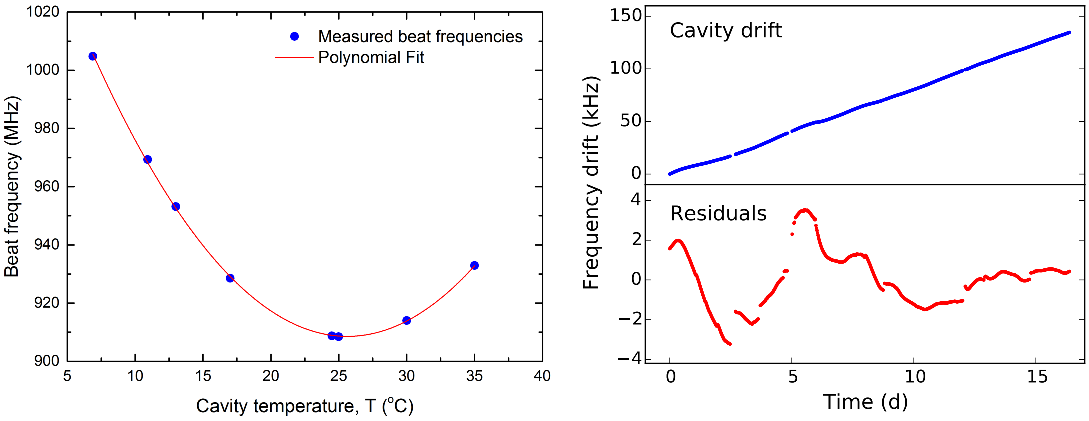

Figure 3.7: The dependence of the clock cavity resonance on temperature, showing a zero-alpha point around 26 °C (left). The long-term drift of our cavity measured against the atoms during two weeks of the ITOC measurement campaign (right).

The next threat to the fixed cavity length is thermal expansion, which dramatically reduces the range of viable materials for cavity construction. Most metals, for instance, have a coefficient of thermal expansion (CTE) in the range of 1×10^{-5}/K, so a typical metal spacer would need unrealistic temperature control below 100 pK for decent frequency stability.

The best spacer materials either possess a zero-crossing in CTE (e.g. Corning ULE, Zerodur, silicon [25, 286]), or more rarely have a CTE which tends to zero as $T \to 0$ K (e.g. sapphire [262]). We have opted to use ULE glass for its convenient zero-crossing near room temperature and because it has about half the mechanical loss of similarly-thermally-stable Zerodur [202]. In order to operate as close as possible to the zero-crossing, we map out the response of our particular ULE spacer in figure 3.7, finding the turning point at around 26 °C. Assuming we suffer from a temperature offset of less than 0.2 K from the zero-crossing we get a residual sensitivity of less than 2×10^{-11}/K, implying that stability at the level of 20 µK over 100 s is sufficient to avoid problematic fluctuations in cavity length.

In typical conditions we observe long-term temperature excursions in the inner shield of a few mK over several days and a fairly slow time constant of around 690 s that ensures that the short-term stability is much better (measured in [46]), though there will on occasion be spikes and settling oscillations that can ruin our prospects for clock measurements for a day or so. The final proof that our temperature control is adequate is the decent frequency stability achieved relative to the Sr clock transition, which shows no threatening features at the relevant timescales (see figure 3.7).

3.1.2.4 Thermal flicker floor: the fundamental limit

Even if we can perfectly control the average cavity temperature, we still get hit by fundamental thermal fluctuations. A fairly small effect is the statistical variance $\langle \delta T^2 \rangle = \frac{k_B T^2}{\rho C V}$ which inevitably arises in any bulk material at thermal equilibrium; this feeds into cavity length fluctuations through the thermo-refractive and thermo-elastic coefficients of the mirror substrate and coating [79]. However, the real killer is Brownian motion, which introduces $1/f$ flicker noise in length fluctuations by (very slightly) exciting the cavity's modes of vibration; in recent generations of optical reference cavity, it is this effect in particular that has limited the local oscillator stability at the most crucial averaging times in the range of $0.1 - 10\,\mathrm{s}$ [99, 101, 124, 135, 166, 202, 251, 279].

It is possible to use finite element modelling of the whole cavity to calculate Brownian noise: utilizing the fluctuation-dissipation theorem, the noise spectrum of cavity length fluctuations can be derived from the (numerically calculable) rate of energy dissipation when oscillating forces are applied to the cavity mirrors [154]. Although these computational models have quite accurately replicated experimental results, they are somewhat outside the scope of this thesis.

Instead we use a simplified model to derive approximate results, splitting the thermal noise into three separate contributions from a cylindrical spacer, a thick mirror substrate, and a thin mirror coating (see figure 3.8) [202]. We consider the displacement $X(t)$ at the spots where the circulating laser beam hits the cavity mirrors, and we are interested in particular in how $X(t)$ responds to small forces $F(f)$ on the mirrors oscillating at different noise frequencies f. We approximate the response using a complex spring constant $(1 + i\phi)k$ defined by:

$$F(f) = -(1 + i\phi)kX(f) \tag{3.6}$$

where k is the spring constant magnitude and ϕ is the 'loss angle' which internally damps oscillations. (ϕ could in theory be frequency dependent, but no dependence has been found for typical cavity materials in the relevant frequency range [202]).

The fluctuation-dissipation theorem implies that there will be thermal excitation of displacements $X(t)$ with a single-sided power spectral density $L_X(f)$:

$$L_X(f) \approx \frac{2k_B T}{\pi f}\left(\frac{\phi}{k}\right) \tag{3.7}$$

where the approximation holds in the limit of small Fourier frequency f compared with any mechanical cavity resonances and in the limit of small loss angle $\phi \ll 1$.

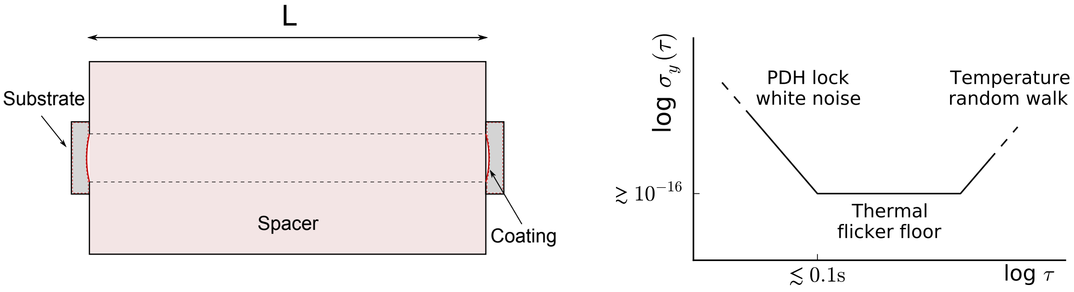

Figure 3.8: The three parts of a reference cavity that contribute roughly separable components to the thermal flicker floor (left), and the typical overall instability curve of a cavity-locked laser (right).

The implications of equation 3.7 for our three-component model are explained in [202], with analytical results given by:

$$L_{\text{spacer}}(f) = \frac{2k_BT}{\pi f}\,\frac{L_0}{3\pi R^2 E_{\text{spacer}}}\,\phi_{\text{spacer}} \tag{3.8}$$

$$L_{\text{substrate}}(f) = \frac{2k_BT}{\pi f}\,\frac{1-\sigma^2}{\sqrt{\pi}E_{\text{substrate}}w_0}\,\phi_{\text{substrate}} \tag{3.9}$$

$$L_{\text{coating}}(f) = \frac{2k_BT}{\pi f}\,2d\,\frac{(1-2\sigma)(1+\sigma)}{\pi E_{\text{substrate}}w_0^2}\,\phi_{\text{coating}} \tag{3.10}$$

where L_0 and R are the length and radius of the cavity spacer, the E_{object} are the Young's modulus, σ is the Poisson's ratio of the substrate, d is the coating thickness, and w_0 is the spot size at the cavity mirror.

The noise contributions are roughly statistically independent, so we can add them up to get a total thermal noise limit to the frequency instability of the cavity-stabilized laser:

$$L_y(f) = \frac{1}{L_0^2}L_X(f) = \frac{1}{L_0^2}\left(L_{spacer}(f) + \sum_{\text{mirrors}}\left(L_{substrate}(f) + L_{coating}(f)\right)\right) \tag{3.11}$$

which we relate to a noise floor in the Allan deviation, $\sigma_y = \sqrt{2\ln 2 L_y(f = 1\,\text{Hz})}$, using one of the identities in table 3.1.

There are a series of implications of the thermal noise eqs. (3.8) to (3.11):

- Smaller substrate loss angle is better. This has motivated a move away from ULE mirror substrates like ours ($\phi = 1.7 \times 10^{-5}$) towards fused silica substrates with very small $\phi = 10^{-6} - 10^{-8}$ [203], enabling room temperature ULE cavities to reach thermal noise floors well below $\sigma_y = 1 \times 10^{-15}$. However, although it might be tempting quickly to swap out the ULE mirrors on our cavity, there are complications due to the mismatch between the thermal expansion of fused silica from the ULE spacer, requiring a tailored ULE ring to compensate [136, 147].

- Longer cavities are better. Since the spacer contribution is usually not dominant the thermal noise floor scales as $1/L_0$, motivating the development of monstrous 39.4 and 48 cm cavities in [251] and [99]. But with increased length, reduction of vibration sensitivity becomes increasingly difficult.

- Smaller coating loss angle is better. Since coating noise is dominant in ULE-fused silica or cryogenic silicon cavities [135], there has been a recent focus on replacing standard dielectric SiO_2/Ta_2O_2 mirror coatings with crystalline $GaAs/Al_{92}Ga_8As$ coatings which can exhibit more than an order of magnitude lower mechanical loss [60] while retaining comparable cavity finesse [61].

- Colder is better. This slightly helps the stability of existing cryogenic silicon cavities at 124 K [135], but is also spurring efforts to build even colder 4 K cubic silicon cavities aimed at low 10^{-17} stability [168].

- Larger spot size w_0 is better, so within the constraints of beam-pointing stability it is best to use of mirrors with a large radius of curvature.

Using our somewhat stubby 7.75 cm room-temperature cavity with ULE substrates and dielectric coatings, we don't really stand a chance of matching state-of-the-art performance. However, the recently developed universal synthesizer is set to save the day.

3.1.3 The universal synthesiser

Following the lead of pioneering work in [2, 100, 120, 191, 195], the universal synthesizer at NPL uses an erbium-doped fibre frequency comb to distribute the stability of one super-stable 1064 nm master laser to a handful of local oscillators operating at diverse wavelengths. Several clocks benefit from the scheme, including 871 nm and 934 nm for Yb^+, 698 nm for Sr, 674 nm for Sr^+, and 9.2 GHz for the Cs primary standard.

From the perspective of the Sr clock, some of the most important features of the universal synthesizer are sketched in figure 3.9. The setup in the universal synthesizer lab is geared towards the generation of a 'transfer beat' at a frequency Δ_{Xfer}

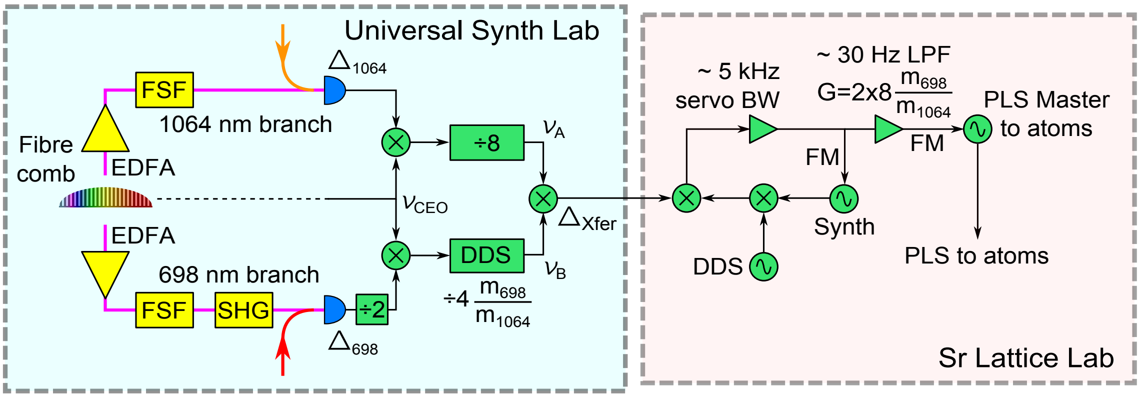

Figure 3.9: Part of the universal synthesizer scheme, showing the electronics used to generate the transfer beat (left) and our method of feeding forward corrections for coherent interrogation of for Sr atoms (right). EDFA: erbium-doped fibre amplifier. FSF: frequency-shifting fibre. SHG: second harmonic generation.

[264], comparing our 698 nm laser with the 1064 nm reference. The idea behind the transfer beat is to provide immunity to comb noise, which is particularly necessary because each tooth of our maser-stabilised comb has quite a broad ∼ MHz linewidth. Propagating the optical beats through all the mixers and dividers in the figure, we find that the noise in comb repetition rate factors out:

$$\nu_A = \frac{1}{8}\left(\nu_{1064} - m_{1064}\nu_{\text{rep}}\right) \tag{3.12}$$

$$\nu_B = \frac{m_{1064}}{8m_{698}}\left(\nu_{698} - m_{698}\nu_{\text{rep}}\right) \tag{3.13}$$

$$\implies \quad \Delta_{\text{Xfer}} = \nu_B - \nu_A = \frac{1}{8}\left(\frac{m_{1064}}{m_{698}}\nu_{698} - \nu_{1064}\right) \tag{3.14}$$

One of the most important concerns for this scheme is whether it can transfer the 1064 nm stability to our laser, i.e. how much excess noise is introduced by the comb? The best published excess instability in comb transfer is 4×10^{-18} at 1 s [195], though this very low noise was based on two frequencies being compared via the same EDFA branch. In a couple of different systems with separate EDFA channels like ours, the differential-mode fibre amplifier noise has been found to introduce significant excess instability at the 10^{-16} level at 1 s [100, 191]. We haven't yet had the tools to carry out a full evaluation in our universal synthesizer system, but the early measurements of transfer beat noise are the basis of some optimism.

When the transfer beat signal was piped through the walls into our lab in late 2015, we needed to think of the best way to use it. The first priority was to analyze

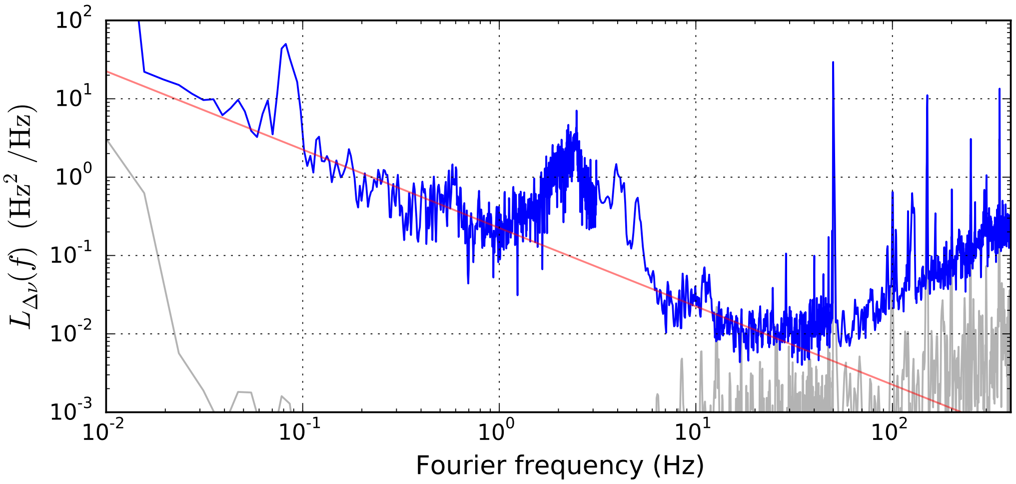

Figure 3.10: Power spectral density of frequency fluctuations of the 698-1064 nm transfer beat, normalised to the Sr clock frequency. The grey line shows the noise floor of the PSD measurement, and the red line indicates flicker frequency noise at $\sigma_y(\tau) = 1.3 \times 10^{-15}$.

its frequency noise spectrum to see what potential it might have to improve our laser stability. To take the spectrum we phase-locked a slave synthesizer to Δ_{Xfer} with a 5 kHz servo bandwidth by actuating on the synthesizer FM input, which was set to a known sensitivity of 100 Hz V^{-1}. The actuation voltage of the servo loop then tracks changes in the frequency Δ_{Xfer} within the servo bandwidth, so all we have to do is send that voltage into a signal analyser to take its fast Fourier transform. The resulting noise spectrum shown in figure 3.10 is the sum of independent contributions from three sources: (1) our 698 nm laser, (2) the 1064 nm reference, and (3) excess comb noise. The spectrum highlights several sources of instability plaguing our lasers:

- The spike at 80 mHz is from 1064 nm crystal temperature controller oscillations.
- The small lump at 0.6 Hz is from vibrations afflicting the 1064 nm AVI.
- The hill around 2 Hz is from optical table oscillations infecting the 698 nm laser.
- The flicker noise corresponding to $\sigma_y \approx 1.3 \times 10^{-15}$ is mostly thermal noise at 698 nm (in good agreement with measurements on the same cavity design [166]).
- The upward turn above 40 Hz is probably phase noise introduced by the comb.

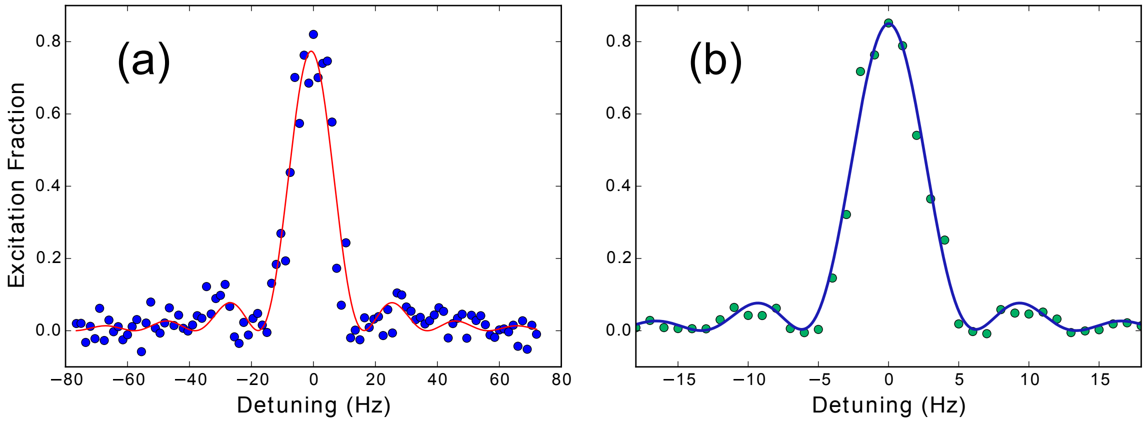

Figure 3.11: Narrow-linewidth Rabi scans of the clock transition in spin-polarised ^{87}Sr. (a) A 50 ms Rabi π-pulse scan without the universal synthesiser and (b) a 150 ms scan with the universal synthesiser.

The dominant sources of noise are from our local 698 nm clock laser, so when we engaged the universal sysnthesiser feedforward we expected to see a significant improvement in the probe coherence. Indeed, this is exactly what happened, as shown in figure 3.11 – where our local laser could only do Rabi pulses up to around 80 ms before suffering reduced contrast, the universal synthesiser could manage coherent pulses of more like 300 ms.

3.2 Stabilising to the atoms

Even by itself, a cavity-stabilised laser acts as an impressively stable clock. However, as the minutes pass by, the cavity will begin to drift in length – whether from thermal fluctuations or from material creep – causing the laser to pull steadily away from the right frequency. So, to achieve the best-performing clock we compensate for this drift by constantly steering the laser onto resonance with the atomic transition. Since the physics of the atom are fixed (or at least we usually assume they are), the atom-locked clock frequency can achieve remarkable stability in the longer term.

3.2.1 Anatomy of an atom lock

In outline, the procedure for locking to a narrow atomic transition goes as follows:

1. Prepare a sample of cold atoms in a well-characterised environment

2. Probe the atoms with light from the local oscillator

3. Measure the fraction F of atoms that were excited by the probe
4. Using the information from this measurement, steer the local oscillator towards resonance

and then repeat 1-4 indefinitely. This indirect method of detecting probe resonance via the atomic excitation fraction is often called 'electron shelving' [66], and it is usually the only viable approach to measure narrow transitions with lifetimes $> 10\,\mathrm{ms}$ $-$ one could imagine alternative detection schemes based on direct absorption or fluorescence from the clock probe beam, but these methods yield too weak a signal due to the very small oscillator strength of the clock transition.

To get the strongest signal, it's best to probe the atoms in a way which is as sensitive to the local oscillator frequency as possible, i.e. at a detuning where the slope $dF/d\bar{\nu}$ is steepest. We also need to avoid any offsets of the locked clock frequency from the atomic resonance, which typically means that we must probe the atoms using a pair of sequences which are symmetrical around resonance $-$ if the sequences have the right symmetry, we can ensure immunity against fluctuations in contrast or exact fringe width. Two common interrogation methods which combine these virtues are depicted in figure 3.12. To lock to one of these spectroscopy features, we repeatedly probe the atoms using one mode of interrogation (high side or $\pi/2$ phase) and then the other mode (low side or $-\pi/2$). After the n^{th} pair of clock cycles we calculate an error signal

$$\epsilon_n = \frac{1}{2}\left(F_{2n} - F_{2n+1}\right) \tag{3.15}$$

which we send into a feedback loop to steer the laser frequency towards atomic resonance. In the feedback loop we correct the laser frequency with the sum of an integrated component $\propto \sum\limits_{n} \epsilon_n$ and a doubly-integrated component $\propto \sum\limits_{m} \left(\sum\limits_{n=1}^{m} \epsilon_n \right)$: the role of the higher-gain single integrator is to capture lock stably to the atomic transition within a few clock cycles ($\sim 10\,\mathrm{s}$ capture time), while the slower second integrator ensures in that any linear drift in the local oscillator does not cause the laser frequency to sag above or below resonance [209].

As it turns out, for reasonably well-tuned loops the long-term instability is not actually affected all that much by the loop gain $-$ rather, the dominant cause of instability is from noise on the error signal ϵ_n.

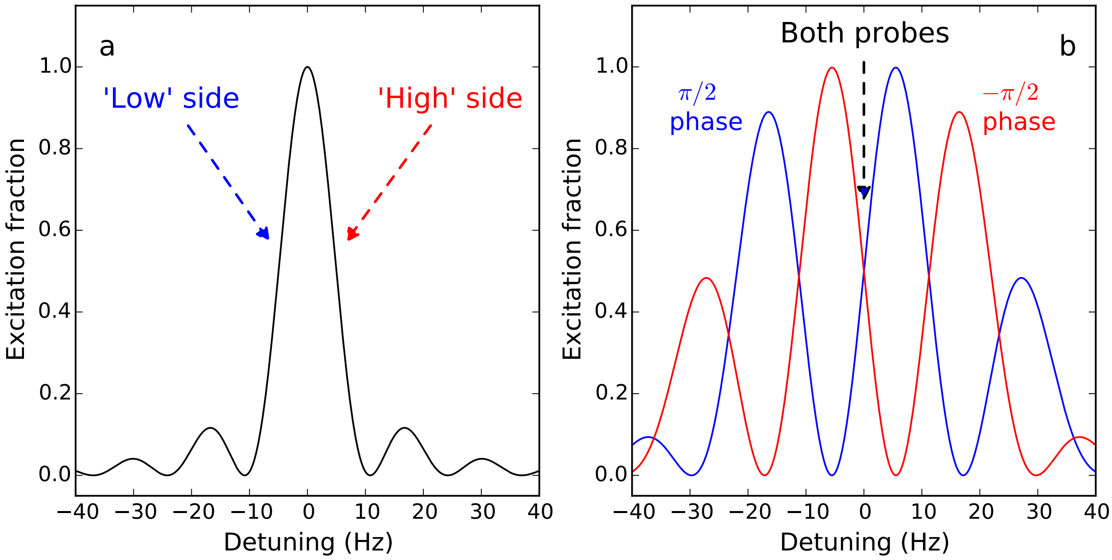

Figure 3.12: Common schemes for locking to an atomic transition. (a) A Rabi π-pulse lock, hopping between the high and the low side of the line to rest at the points of maximum slope. The excitation difference between these interrogation points is unaffected by line-broadening or changes in contrast, thus preventing lock offsets from emerging. (b) A Ramsey lock using a phase step of $\pm\pi/2$ in the first Ramsey pulse to yield a strong discriminant at $\delta = 0$. Again, the crossing point is immune to reductions in fringe contrast.

3.2.2 Clock instability: the error-signal noise limit

We consider a locked clock with noise η_n added to the error signals ϵ_n gathered after each pair of clock cycles. Assuming the local oscillator frequency ν always stays within the linear region on the sides of the fringe, we can write

$$\epsilon_n = D(\bar{\nu}_n - \nu_0) + \eta_n \tag{3.16}$$

where $\bar{\nu}$ is the average clock frequency during the n^{th} pair of cycles, and D is the error-signal discriminant:

$$D := \frac{d\epsilon_n}{d\bar{\nu}_n} \tag{3.17}$$

$$\implies \text{e.g. the Rabi/Ramsey lock :} \qquad D = \frac{1}{2}\left(\frac{dF_{2n}}{d\bar{\nu}_n} - \frac{dF_{2n+1}}{d\bar{\nu}_n}\right) = \frac{dF_{2n}}{d\bar{\nu}_n} \tag{3.18}$$

Now we consider the behaviour of the clock frequency when we close the loop. Assuming a stable feedback loop, in the long term we ought successfully to servo the average error signal to zero:

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \epsilon_n = 0 \tag{3.19}$$

with a rate of convergence set by the integrator (and double integrator) gains.

Substituting this limiting result into eqn 3.16 we end up with a fixed relationship between the error-signal noise η and the frequency error $\bar{\nu} - \nu_0$ when averaged over a large number N of clock cycles:

$$\frac{1}{N} \sum_{n=1}^{N} (\bar{\nu}_n - \nu_0) = \frac{1}{D} \frac{1}{N} \sum_{n=1}^{N} \eta_n \tag{3.20}$$

which we can plug into the definition of the Allan deviation in equation 3.1 to get:

$$\begin{aligned}
\sigma_\nu^2(2NT_c) &= \left\langle \frac{1}{2} \left(\frac{1}{N} \sum_{n=N}^{2N-1} (\bar{\nu}_n - \nu_0) - \frac{1}{N} \sum_{n=0}^{N-1} (\bar{\nu}_n - \nu_0) \right)^2 \right\rangle \\
&\xrightarrow[N\gg 1]{} \frac{1}{D^2} \left\langle \frac{1}{2} \left(\frac{1}{N} \sum_{n=N}^{2N-1} \eta_n - \frac{1}{N} \sum_{n=0}^{N-1} \eta_n \right)^2 \right\rangle \\
&= \frac{1}{D^2} \sigma_\eta^2(2NT_c)
\end{aligned} \tag{3.21}$$

where the factor 2 in $2NT_c$ accounts for the fact that each $\bar{\nu}_n, \eta_n$ covers two clock cycles. So, translating into fractional frequency units we get a long-term instability:

$$\sigma_y^2(\tau) = \frac{1}{\nu_0^2} \sigma_\nu^2(\tau) = \frac{1}{\nu_0^2 D^2} \sigma_\eta^2(\tau) \tag{3.22}$$

where we point out that this same limit holds *regardless of loop gain*, as long as the loop remains stable and the local oscillator stays within the linear region on the side of the fringe.

Equation 3.22 has immediate implications for the design of optical clocks. First, all else being equal, it's best to use a transition with a high resonant frequency ν_0 so that there is more leeway to have some error-signal noise in the system while maintaining good fractional frequency stability. (This is an important advantage of optical over microwave clocks).

Second, the slope at the side of the fringe is generally steeper for narrower fringes $\Delta\nu \sim 1/T$, meaning that for the strongest frequency discriminant D we generally want to extend our probe time T as far as possible. It's no coincidence that the most stable clocks have been those with the highest quality factors $Q := \nu_0/\Delta\nu$ [56, 193].

Finally, we need to minimise the error-signal noise as much as possible. This is the focus of the following sections, where we elaborate on three important sources of noise which typically limit the long-term stability of optical atomic clocks: quantum projection noise [121], detection noise, and the Dick effect [69].

3.3 Sources of instability in the atom lock

3.3.1 Quantum projection noise

Upon projective measurement, each atom collapses into either the ground or the excited state, and no further information about its previous quantum state can be extracted. Thus, if we have N_{at} unentangled atoms per cycle, such that their measurement outcomes are statistically independent, then the measured excitation fraction F in each clock cycle will follow a binomial distribution with a mean F_0 and a variance:

$$(\Delta F)^2_{\text{QPN}} = \frac{F_0(1-F_0)}{N_{at}} \approx \frac{1}{4N_{at}} \tag{3.23}$$

where the approximation holds well if $F_0 \approx 1/2$ [121].

Assuming for now that this is the only noise present on the excitation fraction, we can show how it feeds into an error-signal instability:

$$\begin{aligned}
\sigma_\eta^2(2NT_c) &= \left\langle \frac{1}{2}\left(\frac{1}{N}\sum_{n=N}^{2N-1} \frac{1}{2}(F_{2n} - F_{2n+1}) - \frac{1}{N}\sum_{n=0}^{N-1} \frac{1}{2}(F_{2n} - F_{2n+1}) \right)^2 \right\rangle && (3.24)\\
&= \frac{1}{8N^2}\sum_{m=0}^{4N-1} \left(\langle F_m^2\rangle - \langle F_m\rangle^2\right) && (3.25)\\
&= \frac{1}{2NT_c} T_c (\Delta F)^2_{\text{QPN}} && (3.26)
\end{aligned}$$

where we use that the quantum projection noise is uncorrelated between each cycle and that the mean excitation is the same in every cycle, implying that all the cross-correlation terms follow $\langle F_m F_n\rangle = \langle F_m\rangle^2$.

Applying equation 3.22, we reach a quantum projection noise limit:

$$\sigma_y^2(\tau) = \frac{1}{D^2\nu_0^2}\frac{F_0(1-F_0)}{N_{at}}\frac{T_c}{\tau} \approx \frac{1}{D^2\nu_0^2}\frac{1}{4N_{at}}\frac{T_c}{\tau} \tag{3.27}$$

or, rearranging:

$$\sigma_y(\tau) \approx \frac{\alpha}{Q\sqrt{N_{at}}}\sqrt{\frac{T_c}{\tau}} \tag{3.28}$$

where $Q := \nu_0/\Delta\nu$ is the quality factor, and $\alpha = 1/(2D\Delta\nu)$ is a factor of order unity depending on the exact excitation profile of the interrogation method being used.

These numbers are calculated for ideal Rabi and Ramsey sequences in table 3.2. Note that in practice we usually push for long interrogation times T to try to maximise D, but at some point this also reduces the contrast C of the excitation profile – the effect of this is to decrease the discriminant to $D' \approx D/C$. For further discussion, including the implications for optimum pulse duration, see [209].

3.3.2 Detection noise

Detection noise adds an extra variance $(\Delta F)^2_{\text{det}}$ on the measured excitation fraction in each clock cycle, feeding into clock frequency instability through exactly the same mechanism as quantum projection noise. So, following the same reasoning as in the previous section we get a total instability from detection and quantum projection noise:

$$\sigma_y(\tau) = \frac{1}{D\nu_0}\sqrt{(\Delta F)^2_{\text{det}} + (\Delta F)^2_{\text{QPN}}}\sqrt{\frac{T_c}{\tau}} \tag{3.29}$$

The goal for the experimentalist is (if possible) to make sure that detection noise doesn't significantly affect the clock performance - preferably we'd like to say we're limited by some other more 'fundamental' noise source such as quantum projection or Dick noise. To this end, in the fluorescence detection scheme (see section 2.2.4) we need to pay attention to two sources of detection noise: (1) photon shot noise, and (2) background noise.

To tackle (1), it suffices to make sure that the number of detected photons per ground-state atom is $\gg 1$ so that the atom projection noise is guaranteed to be larger than the photon projection noise. In our case, we estimate a photon detection efficiency of around 4×10^{-3} (limited mostly by viewport diameter 32 mm at distance

	$\Delta\nu$	δ_p	D	α	$\sigma_y(\tau)$
Rabi π-pulse	$0.8/T$	$0.38/T$	$1.9T$	0.69	$8.7 \times 10^{-17}/\sqrt{\tau}$
Ramsey	$0.5/T$	$0.25/T$	πT	$1/\pi$	$5.2 \times 10^{-17}/\sqrt{\tau}$

Table 3.2: Rabi and Ramsey spectroscopy numbers: FWHM linewidth $\Delta\nu$, probe detuning of maximum sensitivity δ_p, discriminant D, inverse signal-to-noise α, and instabilities $\sigma_y(\tau)$ for example parameters $N_{at} = 1000$, $T = 200\,\text{ms}$, $T_c = 800\,\text{ms}$, and $\nu_0 = 429.229\,\text{THz}$.

11 cm from the atoms) and a number of scattered photons per atom of X, meaning that we detect a safe number Y of photons.

The second problem of background detection noise also needs care, though the target noise level depends on several factors (e.g. the number of atoms used, amount of Dick noise). For us it was helpful to use small imaging beams to avoid unnecessary scatter from the viewports, and to point the imaging beams with a significant component along the axis of the atomic beam such that the always-present background beam of atoms was mostly Doppler-shifted far away from probe resonance. Our fancy camera (Andor iXon Ultra) with EM gain also helps, though a decent PMT would be perfectly adequate (and significantly cheaper).

3.3.3 The Dick effect: derivation for a side-of-fringe lock

Ideally, the excitation probability for each clock cycle would be sensitive only to the average frequency error $\bar{\nu}_n - \nu_0$ during that cycle: that way, after a long string of cycles we could correct for the measured $\sum \bar{\nu}_n - \nu_0$ and successfully recover an accumulated clock phase in agreement with the atom's resonant frequency ν_0 (at least at a level limited by detection or quantum projection noise).

Unfortunately, as well as being sensitive to the average frequency of the probe, the excitation probability is *also* sensitive to probe frequency oscillations – during the hundreds of milliseconds of atom capturing, cooling, and preparation, the local oscillator frequency can wander around as it pleases and the atom excitation fraction will be unaffected, whereas if the local oscillator jumps in frequency in the middle of the probe pulse then the excitation will jump too.

As depicted in figure 3.13, this uneven sensitivity results in an aliasing problem. If the local oscillator frequency is fluctuating in tandem with the clock cycle and reaches a peak exactly during the probe pulses, the result will be an offset in the locked clock. If we consider random fluctuations in the local oscillator (as modelled in figure 3.13), over time the fluctuations wander in and out of phase with the clock cycle and so a corresponding wander in the lock offset emerges. The resulting clock instability is known as the Dick effect [69], and it represents the dominant source of frequency instability even in lattice clocks using state-of-the-art, low-noise local oscillators [6, 109, 193, 271].

In this section we derive the Dick instability using the simplified fictional example of a side-of-fringe lock where every clock cycle is the same. The error signal from the n^{th} cycle is then just $\epsilon_n = \delta F_n = F_n - F_0$, where F_0 is the setpoint excitation fraction ($\sim$ half of fringe peak for maximum sensitivity). Using this model system we

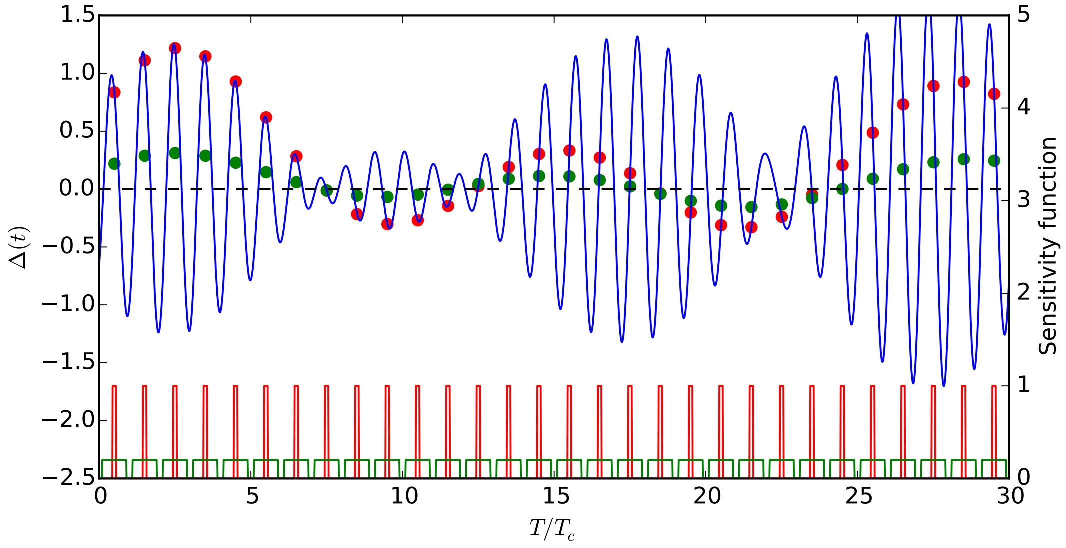

Figure 3.13: Aliasing of random local oscillator noise, focusing in particular on Fourier frequencies filtered at $f = 1/T_c \pm 1/10T_c$ which cause instability at $\sim 10T_c$ timescales. In the red sequence, the very brief sampling of the local oscillator frequency leads to large fluctuations in the lock offset (red points). In the green sequence the local oscillator is sampled over much longer averaging times (green trace), leading to significant smoothing out of the lock offset (green points)

avoid the notational clutter thrown up by interleaved cycles, but we can still reveal the underlying structure of the Dick effect.

3.3.3.1 The single-cycle sensitivity function

If the local oscillator is sitting exactly where it should, then it will excite the equilibrium fraction F_0 of atoms. If the laser deviates from the correct frequency then this excitation fraction will change, resulting in some modified excitation fraction $F = F_0 + \delta F$. Assuming a small time-dependent local oscillator frequency error $\Delta(t) = 2\pi \times (\nu(t) - \nu_0)$ over a single clock cycle $0 \leq t < T_c$, we can define the single-cycle sensitivity function $g(t)$ to fulfil

$$\delta F = \frac{1}{2} \int_0^{T_c} g(t)\Delta(t)dt \tag{3.30}$$

For convenience in later sections, we define $g(t) = 0$ for times outside the first clock sequence, i.e. where $t < 0$ or $t \geq T_c$. The shapes of a couple of typical lattice-clock sensitivity functions are depicted in 3.14: we start with a long period of dead time while a new atom sample is prepared, the sensitivity then climbs to a peak in the

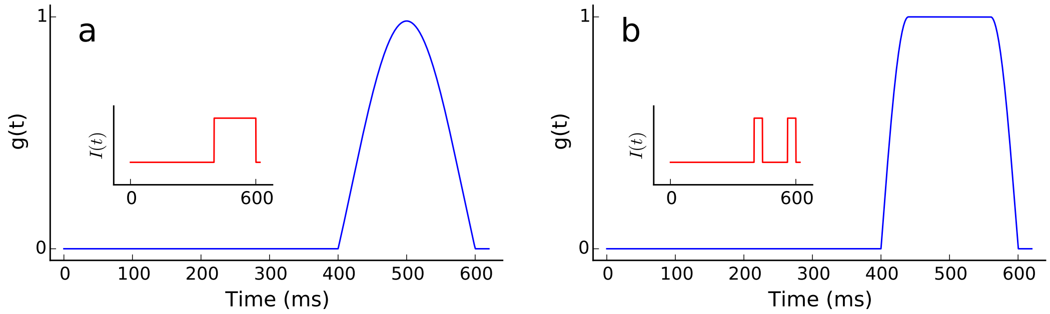

Figure 3.14: Sensitivity functions for typical (a) Rabi and (b) Ramsey sequences. *Insets*: Clock probe intensities over time.

middle of the clock pulse, and finally falls back to zero during the dead time needed for state readout.

In figure 3.14 and throughout the rest of this thesis, we numerically calculate single-cycle sensitivity functions in Python using the derivative equation:

$$g(t) = \lim_{\substack{\epsilon \to 0 \\ \delta t \to 0}} \frac{\left| \langle e| \hat{U}_0(T_c, t+\delta t)\hat{U}_\epsilon(t+\delta t, t)\hat{U}_0(t,0) |g\rangle \right|^2 - \left| \langle e| \hat{U}_0(T_c, 0) |g\rangle \right|^2}{\pi \epsilon \delta t} \tag{3.31}$$

where $\hat{U}_0$ is the atomic evolution operator with the probe laser exactly at the intended frequency, and $\hat{U}_\epsilon$ is the atomic evolution with the probe laser offset from the correct frequency by an extra detuning ϵ. Further treatments of $g(t)$, including derivations of analytical expressions for Rabi and Ramsey spectroscopy, can be found in e.g. [175, 282] – in this thesis, we stick to numerical modelling.

Recalling section 3.2.2, we can express the error signal $\epsilon = \delta F$ as the sum of a signal component and a noise component:

$$\epsilon = D(\bar{\nu} - \nu_0) + \eta \tag{3.32}$$

which allows us to write:

$$\eta = \frac{1}{2}\int_0^{T_c} g_\eta(t)\Delta(t)dt \tag{3.33}$$

$$\text{where} \qquad g_\eta(t) := g(t) - \bar{g} \tag{3.34}$$

The modified sensitivity function $g_\eta(t)$ represents the 'nasty' part of the single-cycle sensitivity function $g(t)$: it adds no useful signal proportional to $\bar{\nu} - \nu_0$, but

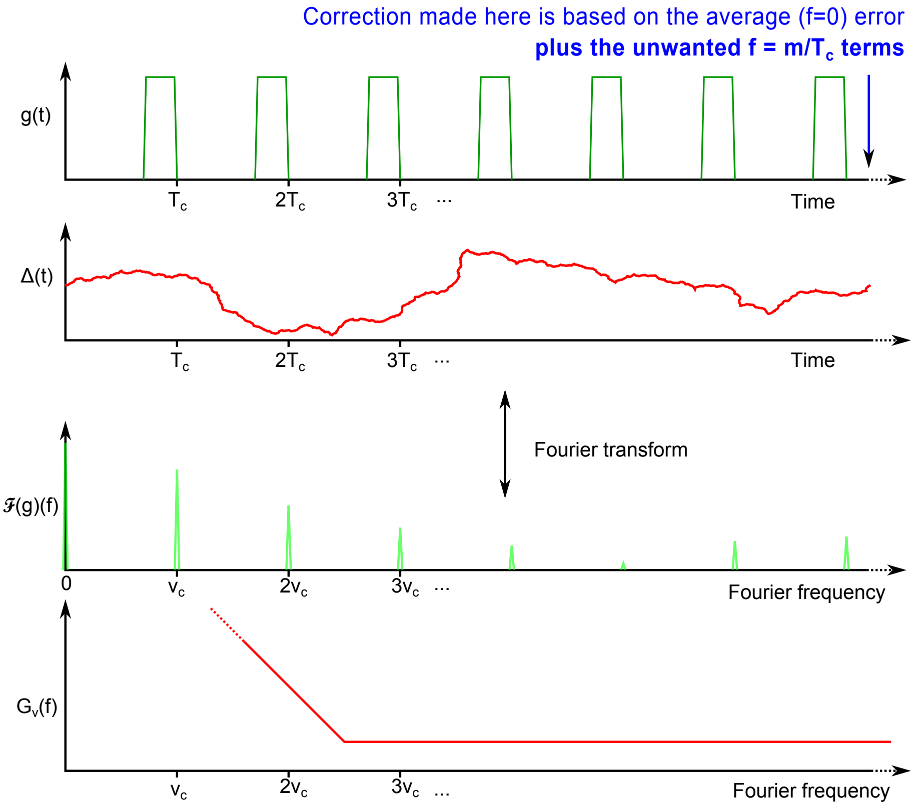

Figure 3.15: To calculate the effect of sampled LO noise, we convert from the time domain instability $\sim |\int g(t)\Delta(t)dt|^2$ to the frequency domain $\sim \int |\hat{g}(f)|^2 L_\nu(f)df$.

only samples noisy local oscillator frequency fluctuations. In the following section we show how this ultimately feeds into clock frequency instability.

3.3.3.2 Error-signal and locked-frequency Dick instabilities

The goal of this section is to derive the instability in error-signal noise η due to sampled local oscillator frequency fluctuations. The rough approach is depicted in figure 3.15: we start in the time domain, considering a long train of clock cycles sampling a noisy local oscillator. Then we transform into the frequency domain, recasting the problem in terms of the local oscillator noise spectrum $S_\nu(f)$ and the Fourier-transformed sensitivity function $\hat{g}(f)$ [1].

We start with the definition of the Allan variance (equation 3.1):

---

[1] We use the Fourier transform notation:

$$\hat{x}(f) = \mathcal{F}\{x\}(f) := \int_{-\infty}^{\infty} x(t)e^{-2\pi i f t}dt$$

$$\sigma_\eta^2(NT_c) = \frac{1}{2}\left\langle \left(\frac{1}{N}\sum_{m=N}^{2N-1} \eta_m - \frac{1}{N}\sum_{m=0}^{N-1} \eta_m \right)^2 \right\rangle \tag{3.35}$$

To make progress with this equation, it is convenient to define an aggregated sensitivity function encompassing both consecutive averaging times:

$$h(t) = \frac{1}{N}\sum_{m=N}^{2N-1} g_\eta(t - mT_c) - \frac{1}{N}\sum_{m=0}^{N-1} g_\eta(t - mT_c) \tag{3.36}$$

which, combined with the relationship between η and $g_\eta(t)$ defined in eqn 3.33, allows us to write:

$$\sigma_\eta^2(NT_c) = \frac{1}{2}\left\langle \left(\frac{1}{2}\int_{-\infty}^{\infty} h(t)\Delta(t)dt \right)^2 \right\rangle \tag{3.37}$$

where we assume for now that no corrections are applied to the local oscillator during the averaging periods, i.e. $\Delta(t)$ is the deviation of the free-running local oscillator from resonance.

Now comes a long string of steps to transform into the frequency domain so that we can relate to a physically measurable quantity, the (two-sided) power spectral density of frequency fluctuations in the local oscillator $S_\nu(f)$:

$$\begin{aligned}
\frac{1}{2}\left\langle \left(\frac{1}{2}\int_{-\infty}^{\infty} h(t)\Delta(t)\ dt \right)^2 \right\rangle &= \frac{1}{8}\left\langle \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} h(t_1)h(t_2)\Delta(t_1)\Delta(t_2)\ dt_1dt_2 \right\rangle \\
&= \frac{1}{8}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} h(t_1)h(t_2)\left\langle \Delta(t_1)\Delta(t_1+(t_2-t_1)) \right\rangle\ dt_1dt_2 \\
\text{Autocorrelation fn } R_\Delta : \quad &= \frac{1}{8}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} h(t_1)h(t_2)R_\Delta(t_2-t_1)\ dt_1dt_2 \\
\text{Convolution } R_\Delta \star h : \quad &= \frac{1}{8}\int_{-\infty}^{\infty} \left(R_\Delta \star h\right)(t)h(t)\ dt \\
\text{Parseval's theorem:} \quad &= \frac{1}{8}\int_{-\infty}^{\infty} \widehat{R_\Delta \star h}(f)\hat{h}^*(f)\ df \\
\text{Convolution theorem:} \quad &= \frac{1}{8}\int_{-\infty}^{\infty} |\hat{h}(f)|^2\hat{R}_\Delta(f)\ df \\
\text{Wiener-Khinchin theorem:} \quad &= \frac{1}{8}\int_{-\infty}^{\infty} |\hat{h}(f)|^2 S_\Delta(f)\ df \\
\text{Angular to normal freq:} \quad &= \frac{(2\pi)^2}{8}\int_{-\infty}^{\infty} |\hat{h}(f)|^2 S_\nu(f)\ df
\end{aligned} \tag{3.38}$$

After all that, the problem reduces to finding the complex Fourier transform $\hat{h}(f)$. To this end, we apply the identity

$$\mathcal{F}\left\{g_\eta(t - mT_c)\right\}(f) = e^{-2\pi i m T_c f}\mathcal{F}\left\{\hat{g}_\eta(t)\right\}(f) \tag{3.39}$$

to every term in equation 3.36, yielding:

$$\begin{aligned}\hat{h}(f) &= \left(\frac{1}{N}\sum_{m=N}^{2N-1} e^{-2\pi i m T_c f} - \frac{1}{N}\sum_{m=0}^{N-1} e^{-2\pi i m T_c f}\right)\hat{g}_\eta(f)\\ &= \frac{1}{N}\left(e^{-2\pi i N T_c f} - 1\right)\sum_{m=0}^{N-1}\left(e^{-2\pi i N T_c f}\right)^m \hat{g}_\eta(f)\end{aligned}$$

which, after summing the geometric series and rearranging some terms, gives:

$$\hat{h}(f) = -2ie^{i\pi(2N+1)T_c f}\frac{\sin^2(\pi N T_c f)}{N\sin(\pi T_c f)}\hat{g}_\eta(f) \tag{3.40}$$

so that:

$$|\hat{h}(f)|^2 = 4\frac{\sin^4(\pi N T_c f)}{N^2\sin^2(\pi T_c f)}|\hat{g}_\eta(f)|^2 \tag{3.41}$$

As depicted in figure 3.16, this expression simplifies in the limit of long averaging times, $N \to \infty$, into a Dirac comb:

$$\lim_{N\to\infty}\frac{2T_c\sin^4(\pi N T_c f)}{N\sin^2(\pi T_c f)} = \sum_{m=-\infty}^{\infty}\delta\left(f - \frac{m}{T_c}\right) \tag{3.42}$$

We reach our final result for the error-signal instability by substituting back into 3.38:

$$\begin{aligned}\sigma_\eta^2(NT_c) &= \frac{(2\pi)^2}{8}\int_{-\infty}^{\infty}|\hat{h}(f)|^2 S_\nu(f)\,df\\ &\xrightarrow[N\gg 1]{} \frac{\pi^2}{NT_c}\sum_{m\neq 0}\left|\hat{g}_\eta\left(\frac{m}{T_c}\right)\right|^2 S_\nu\left(\frac{m}{T_c}\right)\end{aligned} \tag{3.43}$$

for integer m. Note that the $m = 0$ tooth of the comb doesn't contribute because we have set $\bar{g}_\eta = 0$ by definition.

Since we have $\hat{g}_\eta(f) = \hat{g}(f)$ for all $f \neq 0$, we can simplify equation 3.43 by replacing $\hat{g}_\eta(f)$ with $\hat{g}(f)$. If we also convert from error-signal instability to frequency instability (equation 3.22), we find a neat expression for the Dick noise limit:

$$
\begin{aligned}
\lim_{\tau \gg T_c} \sigma_y^2(\tau) = \frac{1}{D^2 \nu_0^2} \sigma_\eta^2(\tau) &= \frac{1}{\tau} \frac{\pi^2}{\nu_0^2 D^2} \sum_{m \neq 0} \left| \hat{g}\left(\frac{m}{T_c}\right) \right|^2 S_\nu\left(\frac{m}{T_c}\right) \\
&= \frac{1}{\tau} \frac{1}{\nu_0^2 |\hat{g}(0)|^2} \sum_{m=1}^{\infty} \left| \hat{g}\left(\frac{m}{T_c}\right) \right|^2 L_\nu\left(\frac{m}{T_c}\right)
\end{aligned}
\tag{3.44}
$$

where in the final step we've re-expressed the discriminant D in terms of the mean sensitivity $\hat{g}(0)$, and we've converted from two- to one-sided power spectral density (from $S_\nu(f)$ to $L_\nu(f)$).

However, there is a caveat to this result: the η instability in 3.43 was derived for a free-running local oscillator - but what happens when we actually lock the clock? Assuming the local oscillator frequency corrections are made at regular intervals it turns out that the η instability is unaffected, and thus equation 3.44 stands as a valid description of the Dick limit of a side-of-fringe lock.

To show this, we need to interpret all the terms contributing to eqn 3.44, converting each tooth of the Dirac comb $\hat{h}(f)$ back into the time domain, i.e. we want to express $h(t)$ as a sum of Fourier components:

$$
h(t) = \sum_{m=-\infty}^{\infty} h_m(t) \tag{3.45}
$$

for some $h_m(t)$ corresponding to the m^{th} comb tooth:

$$
\hat{h}_m(f) = \begin{cases} \hat{h}(f) & \text{for } \frac{m-1/2}{T_c} < f \leq \frac{m+1/2}{T_c} \\ 0 & \text{otherwise} \end{cases} \tag{3.46}
$$

In large N limit, $\hat{h}_m(f)$ is very sharply peaked in the small region around $f = \frac{m}{T_c}$, and is essentially zero everywhere else. Therefore we can make the small angle approximation:

$$
\sin\left(\pi T_c \left(f - \frac{m}{T_c}\right)\right) \approx (-1)^m \pi T_c \left(f - \frac{m}{T_c}\right) \tag{3.47}
$$

So, substituting into eqn 3.40 we get:

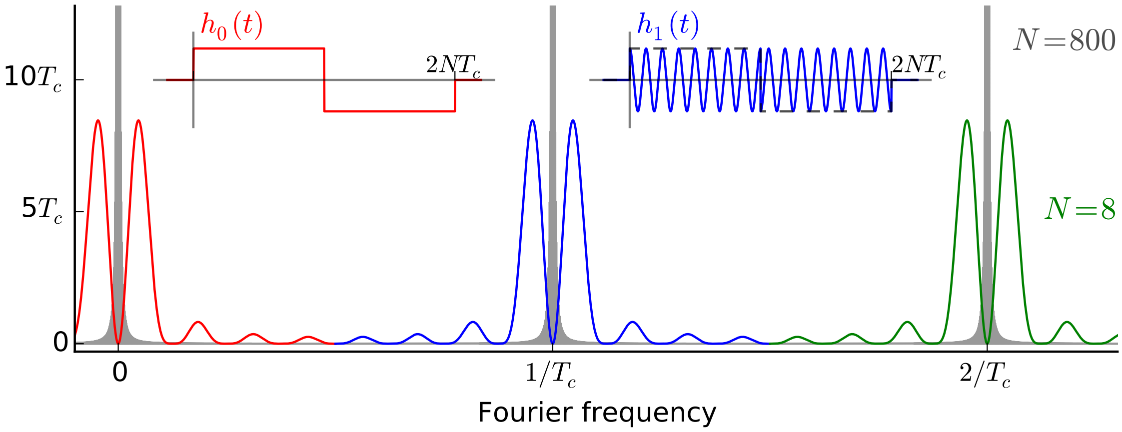

Figure 3.16: Illustration of eqn 3.42 showing how noise at harmonics of the clock cycle are preferentially sampled at long averaging time NT_c. *Insets:* Each tooth of the comb corresponds to an extended Fourier-series basis function $h_m(t)$ (eqn 3.52).

$$\begin{aligned}
\hat{h}_m(f) &= -2ie^{i\pi(2N+1)T_c f}\frac{\sin^2(\pi N T_c f)}{(-1)^m N\pi T_c\left(f-\frac{m}{T_c}\right)}\hat{g}(f) \\
&= -2ie^{i\pi(2N+1)T_c\left(f-\frac{m}{T_c}\right)}\frac{\sin^2\left(\pi N T_c\left(f-\frac{m}{T_c}\right)\right)}{N\pi T_c\left(f-\frac{m}{T_c}\right)}\hat{g}\left(\frac{m}{T_c}\right) \\
&= \left(e^{2i\pi\frac{3NT_c}{2}\left(f-\frac{m}{T_c}\right)} - e^{2i\pi\frac{NT_c}{2}\left(f-\frac{m}{T_c}\right)}\right)\frac{\sin\left(\pi N T_c\left(f-\frac{m}{T_c}\right)\right)}{N\pi T_c\left(f-\frac{m}{T_c}\right)}\hat{g}\left(\frac{m}{T_c}\right) \qquad (3.48)
\end{aligned}$$

where in the first step we've used that $\hat{g}(f) \approx \hat{g}(m/T_c)$ in the non-zero region and that $e^{i\pi(2N+1)m} = (-1)^m$ for integer m, N, and in the second step we've expanded out the sin function into complex exponents.

Finally, to convert this back into the time domain we use the standard Fourier transform results:

$$\mathfrak{F}^{-1}\left(\hat{X}(f)e^{2i\pi af}\right) = X(t-a) \qquad (3.49)$$

$$\mathfrak{F}^{-1}\left(\hat{X}(f-a)\right) = e^{2i\pi at}X(t) \qquad (3.50)$$

$$\mathfrak{F}^{-1}\left(T\frac{\sin(\pi Tf)}{\pi Tf}\right) = \begin{cases} 1 & \text{for } -\frac{T}{2} \le t \le \frac{T}{2} \\ 0 & \text{otherwise} \end{cases} \qquad (3.51)$$

to get:

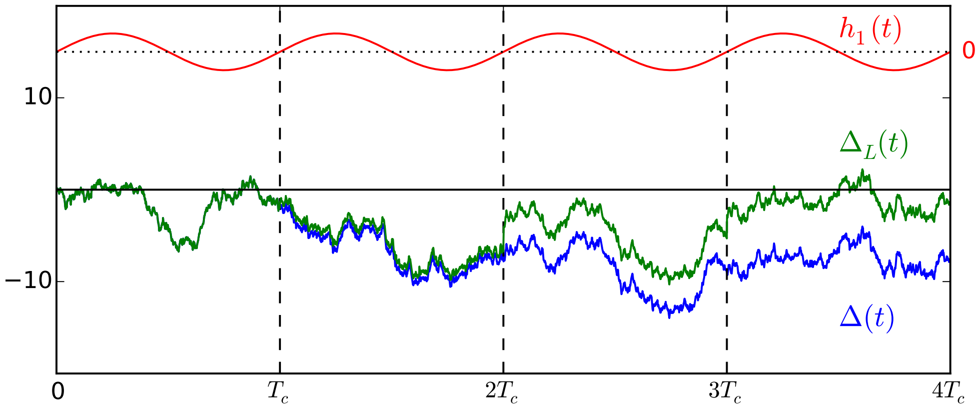

Figure 3.17: A model of a local oscillator with flicker frequency noise, showing atom-locked and free-running frequency traces $\Delta_L(t)$ and $\Delta(t)$. With corrections applied only at the end of each clock cycle, we observe that the h_1-weighted local oscillator frequency fluctuations are the same whether locked or unlocked – the only difference between the traces for any given clock cycle is an overall DC offset. Through the same argument, all h_m-weighted components ($m \geq 1$) that contribute to Dick instability are clearly unaffected by the applied corrections.

$$h_m(t) = \begin{cases} \frac{\hat{g}\left(\frac{m}{T_c}\right)}{NT_c} e^{2i\pi\frac{m}{T_c}t} & \text{for } 0 < t \leq NT_c \\ -\frac{\hat{g}\left(\frac{m}{T_c}\right)}{NT_c} e^{2i\pi\frac{m}{T_c}t} & \text{for } NT_c < t \leq 2NT_c \\ 0 & \text{otherwise} \end{cases} \tag{3.52}$$

which, quite neatly, is just an extension of the Fourier-series basis functions for $g(t)$ over both averaging times (see insets of figure 3.16).

Recalling the previous section, the error-signal instability can be written as a sum over m:

$$\sigma_\eta^2(NT_c) = \frac{(2\pi)^2}{8} \sum_{m\neq 0} \left(\int_{-\infty}^{\infty} |\hat{h}_m(f)|^2 S_\nu(f)\, df \right) \tag{3.53}$$

Keeping in mind this equation, we depict the scenario of a locked local oscillator in figure 3.17: we see that as long as we make our frequency corrections at the same point during each clock cycle, the Dick noise contributions from each integrated $\int h_m(t)\Delta(t)dt$ component remain unchanged. Though such regular corrections aren't the only way of feeding back to the local oscillator frequency, it is hard to see how irregular corrections would decrease the Dick noise unless informed by some external knowledge of the LO fluctuations - so at least for this author, equation 3.44 describes the Dick instability of our locked clock.

3.3.4 The Dick effect: examples

In the previous section we discussed the frequency instability of an atom-locked local oscillator using a simplified side-of-fringe locking scheme. However, in a practical clock we can't use the simple side-of-fringe lock; rather, we need to run the pairs of clock cycles discussed in section 3.2.1. In this section we generalise the Dick effect results to apply to more realistic atomic clock sequences. We find that the Dick instability for a single two-cycle lock is identical to the idealised side-of-fringe lock, but that more care is needed when operating multiple interleaved locks (as required for e.g. characterising systematic frequency shifts). We also explore optimum techniques for systematic evaluation with minimised Dick noise.

3.3.4.1 The single atom lock: derivation

We return to the common case introduced in section 3.2.1 of a single two-cycle lock (e.g. Rabi or Ramsey), where the error signal is calculated after each pair of cycles as $\epsilon_n = \frac{1}{2}(F_{2n} - F_{2n+1})$. To calculate the instability in error-signal noise $\sigma_\eta(\tau)$, we define an error-signal sensitivity function (analogous to equation 3.33) over a pair of cycles $0 \leq t < 2T_c$ to fulfil:

$$\eta = \frac{1}{2}\int_0^{2T_c} g_\eta(t)\Delta(t)dt \tag{3.54}$$

Since the error signal from the first pair of cycles follows $\epsilon = D(\bar{\nu}-\nu_0)+\eta = \frac{1}{2}(F_1-F_2)$, we get:

$$g_\eta(t) = \frac{1}{2}\left((g_1(t) - \bar{g}_1) - (g_2(t-T_c) - \bar{g}_2)\right) \tag{3.55}$$

where $g_1(t)$ and $g_2(t)$ are the single-cycle sensitivity functions for the first and second probe patterns respectively. Since the two pulse patterns are chosen to be symmetric around resonance, they have the opposite sensitivities $g_2(t) = -g_1(t)$. Therefore we can write:

$$g_\eta(t) = \frac{1}{2}\left(g(t) + g(t-T_c)\right) - \bar{g} \tag{3.56}$$

where we've abbreviated using $g(t) = g_1(t)$.

Using this expression for g_η, we can carry out all the steps in section 3.3.3.2 with the only difference being the substitution $T_c \to 2T_c$, finally resulting in a long-term signal instability:

$$\lim_{N\to\infty} \sigma_\eta^2(2NT_c) = \frac{\pi^2}{2NT_c} \sum_{m=1}^{\infty} \left| \hat{g}_\eta \left(\frac{m}{2T_c} \right) \right|^2 L_\nu \left(\frac{m}{2T_c} \right) \tag{3.57}$$

We can simplify by decomposing g_η into its component single-cycle sensitivity functions (Fourier transform of equation 3.55):

$$|\hat{g}_\eta(f)|^2 = \cos^2(2\pi T_c f)\, |\hat{g}(f)|^2 = \begin{cases} \left| \hat{g}\left(\frac{m}{2T_c}\right) \right|^2 & \text{for even m} \neq 0 \\ 0 & \text{for odd m and } m = 0 \end{cases} \tag{3.58}$$

implying that, after all, the only Fourier components that contribute Dick noise are again harmonics of the single cycle frequency $1/T_c$.

Using equation 3.58 and applying equation 3.22 to convert from error-signal instability into the Dick instability of the atom-locked clock frequency, we get the same result as we did for the side-of-fringe lock:

$$\begin{aligned} \lim_{\tau \gg T_c} \sigma_y^2(\tau) = \frac{1}{D^2 \nu_0^2} \sigma_\eta^2(\tau) &= \frac{1}{\tau} \frac{\pi^2}{\nu_0^2 D^2} \sum_{m=1}^{\infty} \left| \hat{g}\left(\frac{m}{T_c} \right) \right|^2 L_\nu \left(\frac{m}{T_c} \right) \\ &= \frac{1}{\tau} \frac{1}{\nu_0^2 |\hat{g}(0)|^2} \sum_{m=1}^{\infty} \left| \hat{g}\left(\frac{m}{T_c} \right) \right|^2 L_\nu \left(\frac{m}{T_c} \right) \end{aligned} \tag{3.59}$$

which is in agreement with the results given in [175, 282].

3.3.4.2 The single atom lock: experimental implications

Since the easy availability of thousands of atoms pushes the quantum projection noise in optical lattice clocks below $10^{-16}/\sqrt{\tau}$ for Rabi pulses of 200 ms or longer, the Dick effect has tended to be dominant even in lattice clocks using state-of-the-art low-noise local oscillators [6, 109, 193, 271]. The effect is compounded by the need to spend a long time in every cycle gathering a new sample of atoms; this dead time not only reduces the average sensitivity $|\hat{g}(0)|^2$, but most importantly it extends the cycle time so that lower frequencies of $1/f$ thermal noise are sampled (see section 3.1).

The effect of dead time is depicted in figure 3.18, where for a fixed Rabi pulse time of 200 ms, we model the Dick effect for dead times of 300 ms and 800 ms. Using

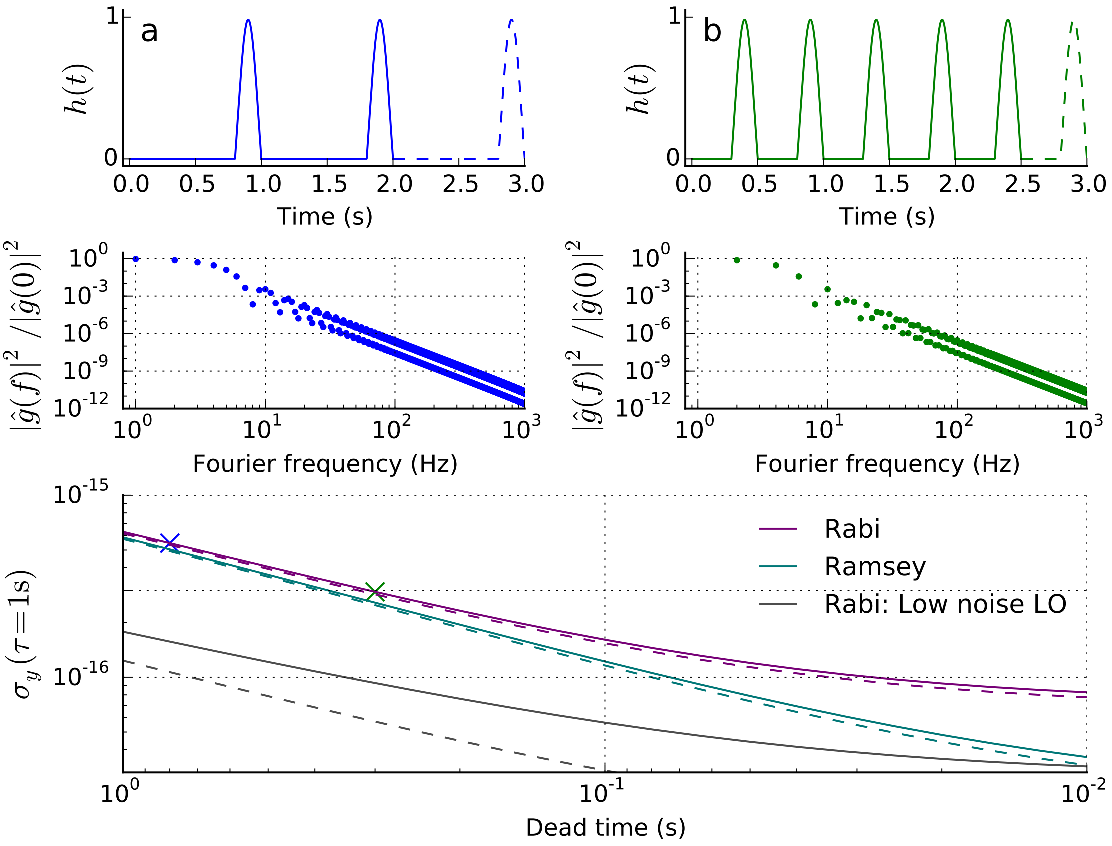

Figure 3.18: Dick noise from clock sequences with different dead times. (a) and (b): the extended sensitivity function $h(t)$ and the resulting spectrum of sampled of local oscillator noise for different Rabi sequences (see text). *Lower plot*: Dick instability in a single Rabi (200 ms pulse) or Ramsey (20 ms pulses and 160 ms dark time) lock. The solid coloured traces have flicker noise 5×10^{-16} and white noise $5 \times 10^{-17}/\sqrt{\tau}$; the grey trace has reduced flicker 1×10^{-16}. The dashed traces have flicker noise only. Crosses mark the Dick instabilities of (a) and (b).

a model local oscillator with a flicker noise floor of $\sigma_y(\tau) = 5 \times 10^{-16}$ summed with a white noise component of $\sigma_y(\tau) = 5 \times 10^{-17}/\sqrt{\tau}$, the Dick instabilities of the two sequences are found to be $2.9 \times 10^{-16}/\sqrt{\tau}$ and $5.4 \times 10^{-16}/\sqrt{\tau}$ for the short and long dead time respectively (see the blue and green crosses on the bottom plot). We can find the cause of the different instabilities by looking at the frequency-domain sensitivity function plots (middle row): at the Fourier frequencies < 10 Hz that contribute the vast majority of Dick noise, the long-dead-time example (blue) samples more frequencies of noise and down to a lower frequency than the short-dead-time example (green).

So, there are clearly some gains to be had in reducing dead time, especially in Ramsey spectroscopy where the sensitivity function is flat-topped (see figure 3.14). This kind of approach has been advocated in [283], where they exploit a non-destructive measurement scheme to avoid having to reload atoms between every clock pulse. However, perhaps the most significant improvements come instead from using a more stable local oscillator − comparing the grey curve with the purple curve, we see that a reduced flicker floor from $\sigma_y(\tau) = 5 \times 10^{-16}$ to $\sigma_y(\tau) = 1 \times 10^{-16}$ leads to a factor of around 3.5 improvement in Dick-limited clock stability. To achieve the same gain in stability by reducing dead time we'd have to shrink a generous 1 s to an unrealistic 70 ms (experimentally we struggle to achieve much less than 400 ms in our Sr system). Although a factor 3.5 in stability may not seem like so much, bear in mind that it saves us a factor 12 in averaging time; spending one day on a clock comparison is a lot less hassle than operating solidly for a 2-week measurement campaign!

3.3.4.3 Multiple atom locks: derivation

In order to evaluate or extrapolate systematic frequency shifts we often need to run several clock servos in an interleaved sequence, with each servo operating under different conditions (e.g. we might run servo 1 with high lattice intensity and servo 2 with low lattice intensity to measure the systematic shift per unit of lattice depth). This mode of operation has two important effects on the Dick instability: (1) the frequency stability of each individual atom lock is degraded due to the extra dead time introduced by all the other locks, and (2) the frequency deviations of the separate servos are correlated to each other, linked by use of same local oscillator. This latter effect has sometimes been overlooked (e.g. in [192, 194]), but we find that it has a significant effect on the instability of difference frequencies between interleaved servos.

We start by considering the frequency stability of each atom lock individually. Unfortunately we can't avoid an escalation of notation. Say we operate k separate servos simultaneously, such that it takes $2k$ clock cycles to gather an error signal for each servo. We can write the error signal for the j^{th} servo in the first set of $2k$ cycles as:

$$\epsilon_j = \frac{1}{2}\left(F_{a_j} - F_{b_j}\right) \tag{3.60}$$

where $a_j, b_j \; \epsilon \; [0, 2k-1]$ index the two cycles used to gather the error signal for the j^{th} servo.

We can then define separate sensitivity functions for the Dick noise η_j on each error signal over times $0 \leq t < 2kT_c$:

$$\eta_j = \frac{1}{2}\int_0^{2kT_c} g_{\eta_j}(t)\Delta(t)dt \tag{3.61}$$

which implies:

$$g_{\eta_j}(t) = \frac{1}{2}\left(g_j(t-a_jT_c) + g_j(t-b_jT_c)\right) - \bar{g}_j \tag{3.62}$$

where g_j denotes the single-cycle sensitivity function for the interrogation method used in cycle a_j. (We assume the pair of cycles a_j, b_j for servo j are symmetrical, exhibiting opposite sensitivity functions – this is the standard operating condition for atom locks).

Following once again the derivation in section 3.3.3.2, but with the replacements $\eta \rightarrow \eta_j$, $g_\eta \rightarrow g_{\eta_j}$ and $T_c \rightarrow 2kT_c$, we reach a final Dick frequency instability in the j^{th} clock servo:

$$\begin{aligned}\lim_{\tau\gg T_c}\sigma^2_{y_j}(\tau) = \frac{1}{D_j^2\nu_0^2}\sigma^2_{\eta_j}(\tau) &= \frac{1}{\tau}\frac{\pi^2}{\nu_0^2 D_j^2}\sum_{m=1}^{\infty}\left|\hat{g}_{\eta_j}\left(\frac{m}{2kT_c}\right)\right|^2 L_\nu\left(\frac{m}{2kT_c}\right) \\ &= \frac{1}{\tau}\frac{1}{\nu_0^2|\hat{g}_{\eta_j}(0)|^2}\sum_{m=1}^{\infty}\left|\hat{g}_{\eta_j}\left(\frac{m}{2kT_c}\right)\right|^2 L_\nu\left(\frac{m}{2kT_c}\right)\end{aligned} \tag{3.63}$$

Before applying this result to tangible examples, we first turn to the next problem of taking frequency sums or differences between servos, for the sake of argument using $\nu_2 - \nu_1$. We can write:

$$\begin{aligned}\sigma^2_{\nu_2-\nu_1}(2kNT_c) &= \left\langle \frac{1}{2}\left(\frac{1}{N}\sum_{n=N}^{2N-1}(\bar{\nu}_{2,n}-\bar{\nu}_{1,n}) - \frac{1}{N}\sum_{n=0}^{N-1}(\bar{\nu}_{2,n}-\bar{\nu}_{1,n})\right)^2\right\rangle \\ &\xrightarrow[N\gg 1]{} \left\langle \frac{1}{2}\left(\frac{1}{N}\sum_{n=N}^{2N-1}\left(\frac{\eta_{2,n}}{D_2}-\frac{\eta_{1,n}}{D_1}\right) - \frac{1}{N}\sum_{n=0}^{N-1}\left(\frac{\eta_{2,n}}{D_2}-\frac{\eta_{1,n}}{D_1}\right)\right)^2\right\rangle\end{aligned} \tag{3.64}$$

where $D_j = d\epsilon_j/d\bar{\nu}_j$ are the discriminants for each lock, and we have converted frequency noise to error-signal noise using equation 3.20. For clarity, the bars over the frequencies $\nu_{j,n}$ now indicate a time-average over the whole of the n^{th} set of $2k$ clock cycles.

Examination of equation 3.64 reveals that we can feed a normalised sensitivity function for the difference in error signal noise $\eta_2 - \eta_1$ into the derivation in section 3.3.3.2:

$$w_{\eta_2-\eta_1}(t) = \frac{\pi}{D_2} g_{\eta_2}(t) - \frac{\pi}{D_1} g_{\eta_1}(t) \tag{3.65}$$

where each g_{η_j} is given by equation 3.62 and we insert a π to account for the relation $D_j = \pi \bar{g}_j$). This yields a final result for the Dick instability in the difference frequency $y_2 - y_1$:

$$\lim_{\tau \gg T_c} \sigma^2_{y_2-y_1}(\tau) = \frac{1}{\tau}\frac{1}{\nu_0^2} \sum_{m=1}^{\infty} \left| \hat{w}_{\eta_2-\eta_1}\left(\frac{m}{2kT_c}\right)\right|^2 L_\nu\left(\frac{m}{2kT_c}\right) \tag{3.66}$$

where the unfortunate shift in notation from unnormalised sensitivity g to the normalised sensitivity w is necessary to treat cases where the different servos j have different discriminants D_j.

It is clear that this result can be extended by analogy to find Dick instabilities for other sums or differences between servo frequencies: if we are interested in some combination of frequencies $\sum_j c_j \nu_j$ for any coefficients c_j, we simply need to substitute into 3.66 the appropriate normalised sensitivity function $w_{\sum_j c_j \eta_j} = \sum_j c_j \pi g_{\eta_j}/D_j$.

3.3.4.4 Multiple atom locks: experimental implications

We consider an example set of sequences which we have run quite often in our lab: four interleaved servos with ^{87}Sr. We depict this in figure 3.19. In order to evaluate a systematic shift between two regimes 'A' and 'B' (e.g. high/low lattice intensity, high/low density, high/low atom temperature) we stabilise 4 frequencies, ν_{A+}, ν_{A-}, ν_{B+}, and ν_{B-}, where the '+' and '−' subscripts denote locks to a spin-polarised sample of atoms in the $M_F = +9/2$ and $M_F = -9/2$ states respectively. This enables us to calculate two averaged frequencies:

$$\nu_A = \frac{1}{2}\left(\nu_{A+} + \nu_{A-}\right) \tag{3.67}$$

$$\nu_B = \frac{1}{2}\left(\nu_{B+} + \nu_{B-}\right) \tag{3.68}$$

for which the linear Zeeman shift cancels out (see section 4.1.2).

There are two questions we are concerned with:

1. How stable is each average frequency ν_A, ν_B by itself? This is important if we want to evaluate our clock during a measurement campaign: we can operate ν_A as our 'reference' frequency for comparison with other atomic clocks, while simultaneously monitoring the systematic shift $\nu_A - \nu_B$.

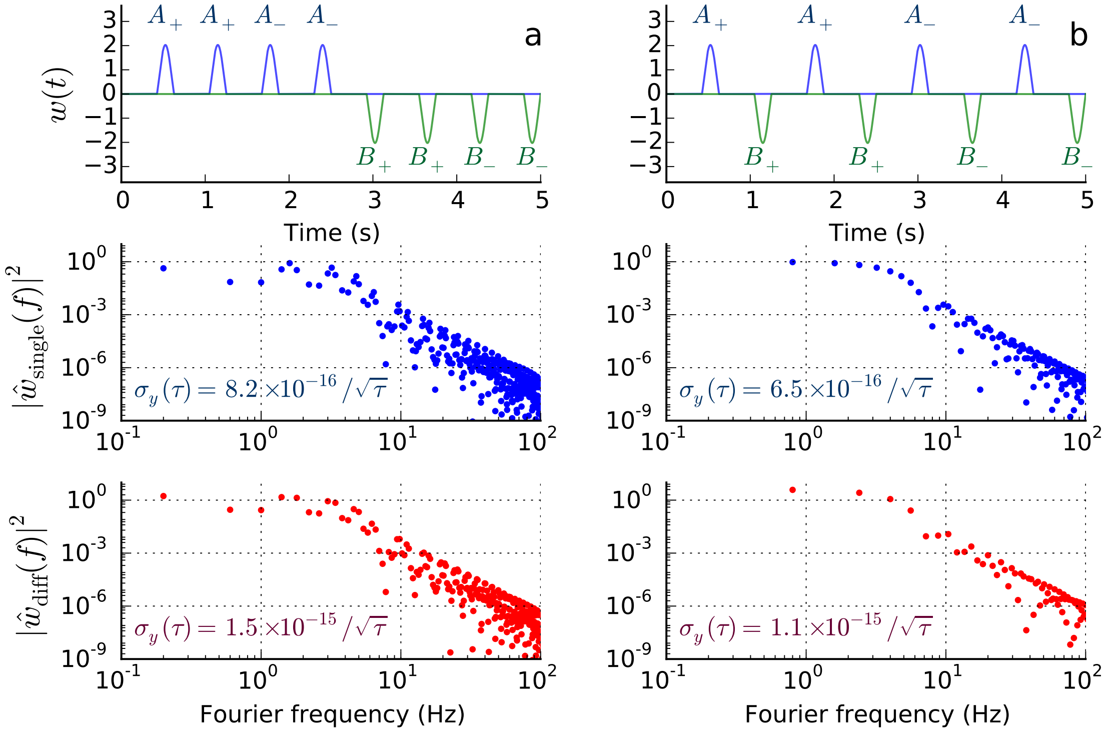

Figure 3.19: Dick noise in interleaved atom locks (see text). *Top row:* Normalised sensitivities $w(t)$ for ν_A and ν_B. *Middle row:* Sampled noise for the average ν_A. *Bottom row:* Sampled noise for the difference $\nu_A - \nu_B$. The inset numbers are the Dick instability for flicker 5×10^{-16} and white noise $5 \times 10^{-17}/\sqrt{\tau}$.

2. How stable is the frequency difference $\nu_A - \nu_B$? This tells us how quickly we can measure the systematic frequency shift between the different operating conditions.

Answers are depicted in figure 3.19 for two different orderings of the eight clock cycles. The ordering turns out to be quite important – in particular the straightforward approach, running all 4 cycles for ν_A then all 4 for ν_B, leads to worse stability than interleaving the cycles one by one. The reason for this can be seen in the Fourier frequency plots: the $4+4$ sequence samples lower frequencies of local oscillator noise than the $1+1$ sequence, making it more vulnerable to the worse end of the $1/f$ tail of local oscillator noise.

Figure 3.19 also reveals a further important point: the correlations between ν_A and ν_B have a significant impact on the instability of the difference frequency $\nu_A - \nu_B$. Unfortunately, the impact is rather unhelpful for accurate frequency measurements, with correlated instability larger than the uncorrelated value:

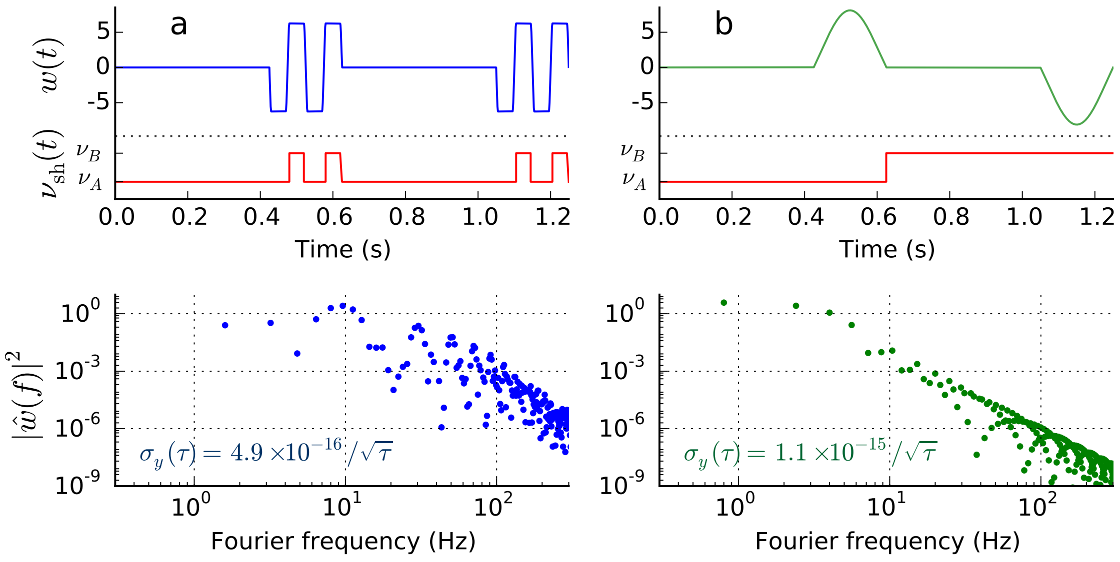

Figure 3.20: Dick noise in different methods of systematic shift evaluation (see text). (a) Normalised sensitivity functions $w(t)$ and the resulting sampled noise components for a 3-spin-echo sequence. (b) The same information for interleaved Rabi sequences. The inset numbers are Dick instabilities for flicker floor 5×10^{-16} and white noise $5\times10^{-17}/\sqrt{\tau}$.

$$\sigma_{y_A - y_B} = 1.1\times10^{-15}/\sqrt{\tau} > 9.2\times10^{-16}/\sqrt{\tau} = \sqrt{\sigma_{y_A}^2 + \sigma_{y_B}^2} \tag{3.69}$$

3.3.4.5 Spin-echo pulses: faster systematic evaluation

The standard method of evaluating a systematic shift, at least if you only have one physics package at hand, is to carry out the type of sequence discussed in the previous section: interleaving two sets of Rabi cycles, one in regime A and one in regime B, we implement separate servos to the different sets of cycles and monitor the frequency difference $\nu_A - \nu_B$. For the sake of argument we consider here a single cycle in regime A corresponding to low lattice intensity, followed by a single cycle in regime B corresponding to high lattice intensity, and we contemplate the instability in resulting measured lattice shift $\nu_A - \nu_B$. The Dick noise in this scenario is depicted in column (b) of plots in figure 3.20.

An alternative method of evaluating systematic shifts in a single physics package is to use spin-echo sequences like the one depicted in the column (a) of plots. Here, the two regimes are interleaved *during* each interrogation pulse. Since it's a spin echo sequence, the excitation fraction at the end of the spin-echo pulse is to first order insensitive to the overall detuning of the interrogating laser. Instead, by interleaving the two regimes at the midpoints of each echo pulse we can make the excitation fraction directly dependent on the systematic frequency shift $\nu_A - \nu_B$. The advantage of this approach over the separate Rabi sequences is evident in the bottom plots: the frequency difference $\nu_A - \nu_B$ samples higher-frequency components of noise from the local oscillator, thus making it considerably less vulnerable to the flicker frequency noise.

In practice, the spin-echo scheme isn't easily applicable to any of the frequency shifts plaguing the optical lattice clock, primarily because it's hard to change any of the relevant shifts quickly enough during the pulse while avoiding undesired effects. For instance the BBR and collisional shifts cannot be switched mid-pulse, the lattice intensity can't be ramped too suddenly or it would heat the atoms too much, and the magnetic field is difficult to change quickly due to eddy currents in our chamber. Nice idea though – maybe someone will find a use for it.

3.4 How stable could lattice clocks get?

With the systematic uncertainty of optical lattice clocks heading towards the sub-10^{-18} level, we'd ideally like to see further progress in clock stability too – even the most stable lattice clock takes a tiresome 28 hours to measure with 5×10^{-19} statistical uncertainty, or a solid month to reach 1×10^{-19} [6]. It would be more practical to exploit the exceptional accuracy of lattice clocks if the required measurement time could be reduced from days to minutes. Based on our models of the Dick effect in the previous section, there are two possible approaches to reaching instability below $1 \times 10^{-16}/\sqrt{\tau}$: either (1) build even better local oscillators, or (2) explore Dick-free methods of interrogating the atoms.

Much of the recent progress in stability in lattice clocks has been based on the first approach, exploiting new innovations in the design of ultrastable cavities; the most stable lattice clock yet built [6] wins primarily by exploiting the most stable cavity yet built with an exceptionally low thermal noise floor of $\sigma_y = 8 \times 10^{-17}$ [99]. However, as we saw in section 3.1 cavity-stabilised local oscillators are already quite a mature technology, knocking simultaneously against the limits of thermal noise and vibration sensitivity – further advances are getting increasingly expensive and experimentally challenging. Promising alternatives to cavities are being pursued, based on spectral hole-burning in cryogenic Eu^{3+}:$Y_2Al_5O_{12}$ crystals [266], superradiant lasers [35, 36, 180], or non-linear cavity-enhanced spectroscopy of narrow transitions [177, 285], but these approaches are some distance away from matching the best cavities – which is a little depressing given that a cavity is basically just a heavily engineered box with a couple of mirrors stuck to it.

It is tempting, then, to explore the second route to more stable clocks: Dick-free interrogation of atoms. In practice this would require at least two physics packages: while we stop to detect, load, and prepare new atoms in one physics package, we could still be probing the atoms in the other; the composite sensitivity function can then be made constant, thereby eliminating the Dick effect completely (see equation 3.59). A couple of possible sequences using two physics packages are depicted in figure 3.21. Stability improvements have already been demonstrated in a microwave clock [26] using the reasonably flat sensitivity of a standard Ramsey sequence applied to two atomic samples. However, maximum immunity to the Dick effect needs specially shaped pulses. One option, plotted in column (b) of the figure, is to use time-dependent Rabi frequencies in the first and last pulse respectively of:

$$\Omega_{\text{rise}}(t) = \frac{A}{t_p}\frac{\sin\left(\frac{\pi t}{t_p}\right)}{\sqrt{1-\sin^4\left(\frac{\pi t}{2t_p}\right)}} \qquad \Omega_{\text{fall}}(t) = \frac{A}{t_p}\frac{\sin\left(\frac{\pi t}{t_p}\right)}{\sqrt{1-\cos^4\left(\frac{\pi t}{2t_p}\right)}} \tag{3.70}$$

where the clock error signal is assumed to be derived using $\pm\pi/2$ phase jumps in the first pulse (i.e. not by jumping the probe detuning – the two approaches lead to slightly different sensitivities).

The big obstacle to Dick-free interrogation isn't the pulse shaping, of course, but the complication of having to build multiple physics packages; it takes long enough to build just the one! And then in order to measure the stability of a Dick-free clock properly, we'd probably need a second Dick-free clock to compare with – at which point we need not just two, but four physics packages.

Thankfully it is sometimes possible to compare two clocks without being affected by Dick instability, even if each clock by itself is still vulnerable to the Dick effect. This can be achieved by interrogating the two atomic samples at the same time and using the same local oscillator, so that the clocks sample exactly the same laser noise. In this 'synchronous' interrogation the Dick noise of each clock cancels out in common mode, allowing dramatic improvement in the precision of the clock comparison for a given averaging time [132, 261]. Synchronous interrogation is especially useful for fast evaluation of systematic shifts, as demonstrated when mapping out the BBR shift between two cryogenic Sr optical lattice clocks in [271]. In combination with methods of correlation spectroscopy [58], synchronous interrogation could also be used to resolve

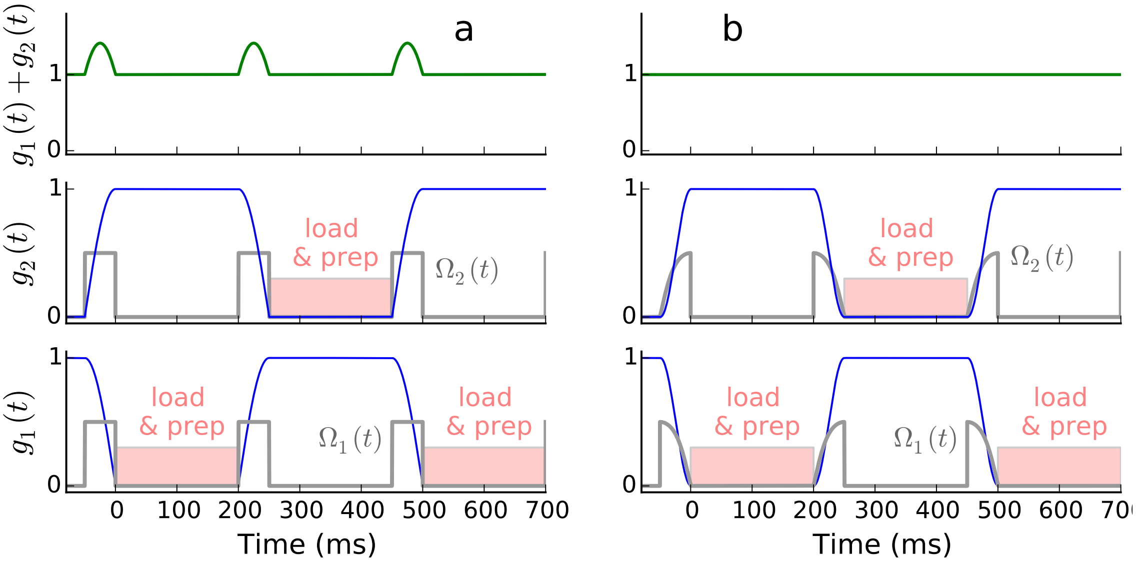

Figure 3.21: Zero-dead time clocks employing (a) standard Ramsey or (b) shaped-pulse Ramsey sequences (as defined in equation 3.70) in two physics packages.

ultra-narrow high-contrast features with interrogation pulses much longer than the local oscillator coherence time, enabling higher Q and thus lower quantum projection noise [55]. However, when trying to make measurements between independent systems or between clocks operating at significantly different frequencies, synchronous interrogation is often not an option; in such cases a Dick-free measurement may only possible via the zero-dead-time approach.

Assuming we do build an optical lattice clock free of the Dick effect, what's next? The limit would then be quantum projection noise (see section 3.3.1), which gives a clock stability of around $3 \times 10^{-17}/\sqrt{\tau}$ for a zero-dead-time clock with 300 ms Ramsey pulses and 1000 atoms. Following through on a frequency measurement with such a clock, it would take an hour to reach 5×10^{-19} precision, or a day to reach 1×10^{-19}. Can we do better? We could try adding more atoms, but this only yields a $\sqrt{N}$ enhancement in stability while introducing collisional effects that can significantly reduce contrast in narrow clock features [175]. Another approach is to extend the Ramsey pulse time to achieve higher Q: Although this would typically reduce contrast because of local oscillator decoherence [209], it may be possible to extend the coherent interrogation time using weak non-demolition measurement (the 'atomic phase lock') [240] or, rather laboriously, by stacking lattice clocks together in a series of stabilization stages with increasing Q [40].

However, perhaps the most scalable approach to reducing quantum projection noise is to introduce entanglement between the atoms, with the aim of squeezing the variance in excitation fraction at the end of the clock pulse. In a 'spin-squeezed' clock, the quantum projection noise could theoretically be reduced by up to an extra factor $\sqrt{N}$ [138, 287], making more accessible a clock instability below the level of $1 \times 10^{-17}/\sqrt{\tau}$. Several methods have been proposed to apply a spin-squeezing Hamiltonian in lattice clocks, e.g. Rydberg dressing [88], non-demolition measurement [142, 181], or transport of an auxiliary entangling atom [281], though none have yet been successfully implemented. When such a system is finally realised it could be used to make measurements with unprecedented precision and speed, aiding the searches for topological dark matter [68] and position- and time-variation of fundamental constants [31, 91, 116, 272].

Chapter 4

Clock accuracy

The accuracy of an atomic clock is determined by how well we can characterise all the systematic shifts to its frequency. The clock-maker therefore has to be meticulous, precisely measuring every possible systematic shift to make its residual uncertainty as small as possible. Once this process is complete, all the shifts are compiled into an error budget (see e.g. table 4.1) so that the clock's total systematic uncertainty can be calculated. As mentioned in chapter 1, the total systematic uncertainty equates to the expected relative inaccuracy of the clock – it tells us how much timing error we'd expect the clock to build up in the long term relative to an ideal clock of the same species.

In this chapter we begin with an analysis of the systematic shifts in ^{87}Sr, the more conventional isotope for building optical lattice clocks. Next, we focus on the less well-characterised ^{88}Sr clock, introducing the new methods of modified hyper-Ramsey and multi-photon spectroscopy. Finally, we describe how lattice clocks with either isotope have the realistic potential to measure time at the 19th decimal place.

4.1 Systematic evaluation of the ^{87}Sr clock

The overall results of the systematic evaluation of our ^{87}Sr lattice clock are shown in table 4.1. We estimate a total fractional systematic uncertainty of around 4×10^{-17}, although the DC Stark shift still awaits proper characterisation (see section 4.1.4). It is important to bear in mind, however, that this total uncertainty is only our best estimate of the ^{87}Sr clock's expected inaccuracy – the table accounts for the known unknowns, but we can't account for any 'unknown unknowns' which may be lurking undiagnosed in our experiment. In order to prove that our clock really is as accurate as we think, we will ultimately need to compare directly against another clock of comparable or better accuracy. For now, we do have some reassuring preliminary

Table 4.1: Systematic shifts and error budget for the ^{87}Sr lattice clock

Systematic effect	Correction ($\times 10^{17}$)	Uncertainty ($\times 10^{17}$)
BBR chamber	474.6	1.0
BBR oven	0.1	0.1
Lattice shifts	0	1.0
Probe Stark	0	0.3
DC Stark	?	?
2nd-order Zeeman	46.0	0.3
Density	0.8	0.8
Line-pulling	0	2.8
1st-order Doppler	0	2.9
Background gas collisions	0	0.2
Servo error	0	0.2
Total	521.5	4.4

evidence that our clock is ticking at roughly the right rate: a measurement against UTC in June 2015 reveals an absolute frequency of 429 228 004 229 873.06(15) Hz after all known systematics are compensated for, in excellent agreement with other reported ^{87}Sr clock measurements around the world [92].

Many of the systematic shifts in table 4.1 are evaluated using the technique described in section 3.3.4 of *interleaved self-comparison*. According to this technique, clock servo 'A' is interleaved with clock servo 'B' , where'A' and 'B' operate under different conditions (e.g. high/low lattice intensity, high/low density, high/low atom temperature). The servo frequency difference $\nu_B - \nu_A$ is recorded and matched against a theoretical model of the shift, enabling us to extrapolate an implied systematic frequency shift under 'ordinary' operating conditions.

However, our estimate of each systematic shift always has some uncertainty. For example, in the method of interleaved self-comparison, uncertainty is inevitably introduced by noise in the data. On the other hand, large uncertainties can sometimes arise from our imperfect theoretical knowledge of the Sr atomic structure. In this section, we run through all the systematics in table 4.1, we outline our methods of evaluation, and we discuss prospects for reduction in their uncertainties.

4.1.1 Blackbody radiation

The systematic shift which typically introduces the largest uncertainty into room-temperature Sr lattice clocks is from the black-body radiation (BBR) bathing the atoms [81, 160, 193]. Indeed, BBR is perhaps the biggest weakness of Sr compared with other types of lattice clock: the fractional shift is 2×, 15×, and 35× larger in Sr than in Yb, Mg, or Hg respectively. Because of this, there has been a concerted effort in recent years to characterise the BBR shift to the highest possible accuracy in Sr [184, 185, 193, 230], and also in Yb, the next-most-sensitive species [21, 239].

4.1.1.1 BBR shift: theory

As explained in e.g. [106, 185], the BBR electric field is seen as approximately 'static' from the perspective of the atoms − most of the black-body radiation oscillates at frequencies much slower than any transitions from either clock state. In the static approximation, the BBR shift $\Delta\nu_{\mathrm{st}}$ is given by

$$h\Delta\nu_{\mathrm{st}} = -\frac{1}{2}\Delta\alpha(0)\left\langle|E|^2\right\rangle_T \tag{4.1}$$

where $\Delta\alpha(0) = \alpha_e(0) - \alpha_g(0)$ is the net DC polarisability of the clock transition, and where the time-averaged mean-square BBR electric field is given by

$$\left\langle|E|^2\right\rangle_T = \frac{8\pi^5 k_b^4}{15\epsilon_0 h^3 c^3}\cdot T^4 \tag{4.2}$$

or, in more user-friendly units, $\left\langle|E|^2\right\rangle_T = (831.94\,\mathrm{V/m})^2 \cdot (T/300\,\mathrm{K})^4$.

$M1$ and higher-order multipole BBR shifts also affect the Sr atom, but it turns out they can be neglected: the $M1$ correction is just 24 µHz for the $^3\mathrm{P}_0$ state at 300 K [214], while no $M1$-allowed transitions exist from the ground state. In principle, tensor shifts from the static $E1$ interaction should also be taken into account in inhomogeneous BBR environments [82], but these are considerably suppressed in lattice clocks due to the lack of electronic angular momentum in either clock state.

In addition to the 'static' shift, there is also a 'dynamic' correction due to the non-zero frequency of the BBR radiation. In Sr, the strong $5s5p^3\mathrm{P}_0$ - $5s4d^3\mathrm{D}_1$ transition at 2.6 µm exacerbates the dynamic correction due to an unfortunate proximity to the peak BBR wavelength, which is around 10 µm at 300 K. Including all the relevant powers of T for $T \lesssim 300\,\mathrm{K}$, the total BBR shift can be written in the form:

Table 4.2: Coefficients in eqn 4.3 for the BBR shift in atomic Sr [185, 193, 230]

	β_{st}	$\beta_{\mathrm{st}}\eta_1$	$\beta_{\mathrm{st}}\eta_2$	$\beta_{\mathrm{st}}\eta_3$	Total $\Delta\nu_{\mathrm{dyn}}$
Shift coefficient (Hz)	-2.13023	-0.1291	-0.01616	-0.00343	-0.1486
Uncertainty (Hz)	0.00006	0.0006	0.00007	0.00002	0.0007
Shift coefficient ($\times 10^{-18}$)	-4962.93	-300.7	-37.6	-7.97	-346.3
Uncertainty ($\times 10^{-18}$)	0.14	1.4	0.2	0.03	1.6

$$\Delta\nu_{\mathrm{BBR}} = \beta_{\mathrm{st}} \left(\frac{T}{T_0}\right)^4 + \beta_{\mathrm{st}} \left[\eta_1 \left(\frac{T}{T_0}\right)^6 + \eta_2 \left(\frac{T}{T_0}\right)^8 + \eta_3 \left(\frac{T}{T_0}\right)^{10} + \mathcal{O}\left(\frac{T}{T_0}\right)^{12}\right] \tag{4.3}$$

where $T_0 = 300\,\mathrm{K}$, and where the static shift coefficient β_{st} and dynamic shift coefficients η_i are given in table 4.2[1]

4.1.1.2 BBR shift: experiment

For a given Sr lattice clock, the total BBR shift has two distinct uncertainties associated with it: (1) The theoretical uncertainty in the BBR coefficients listed in table 4.2, and (2) the experimental uncertainty in the absolute temperature of the BBR environment.

Our lattice clock, like most Sr clocks operated at room temperature [81, 160], is limited by experimental uncertainty in the absolute temperature of the BBR environment. At 300 K, the fractional frequency sensitivity to small changes in temperature is 7×10^{-17}/K, implying that the absolute temperature of the BBR bath would have to be characterised to within 20 mK in order to achieve a coefficient-limited uncertainty of 1.6×10^{-18}. Measuring the BBR at this level is a substantial challenge, especially given various unavoidable heat sources – most prominently the Sr oven and the MOT coils – surrounding the chamber.

---

[1] Since the numbers in table 4.2 are not directly sourced, I'll explain their origins here. The static BBR coefficient and its uncertainty are taken directly from [185]. The dynamic BBR coefficients are extracted from table VI in [230], except for a scaling of the $5s5p^3\mathrm{P}_0$ - $5s4d^3\mathrm{D}_1$ contribution by a factor 2.172/2.180 to account for a later, improved measurement of the D-manifold lifetime in [193]. The uncertainties in the dynamic shift coefficients arise from the 0.46% uncertainty in that lifetime measurement: according to [230], the next-biggest uncertainty – i.e. the theoretical uncertainty in $^3\mathrm{D}_1$ - $^3\mathrm{P}_0$ branching ratio – should only contribute below the 0.3% level.

We observe that $\beta_{\mathrm{st}}\eta_2$ and $\beta_{\mathrm{st}}\eta_3$ *aren't* negligible at the 10^{-18} level, in contrast with the assumption apparently made in some landmark Sr clock papers [193, 271] where $\mathcal{O}(T/T_0)^8$ terms were discarded!

In our experiment, we monitor BBR using eleven Pt-100 temperature sensors glued externally to the science chamber. The sensors are calibrated to an absolute accuracy of ±10 mK, but in practice the dominant source of uncertainty is actually the temperature inhomogeneity over the chamber, which ranges from 300 to 900 mK pk-pk depending on random fluctuations in operating conditions. This inhomogeneity is mostly caused by the temperature mismatch between the $\sim 15\,°\mathrm{C}$ water-cooled MOT coils and the rest of the room at $\sim 21\,°\mathrm{C}$, though an additional source of temperature gradients is the heating effect of a few °C as the cooling water propagates from the outer to the inner windings of the MOT coils.

To derive a net frequency shift from the inhomogeneous BBR environment, our approach is similar in principle to the line-of-sight ray-tracing model described in [174]: we numerically model the total BBR field at the atoms as a solid-angle- and emissivity-weighted sum of contributions from the various components of the science chamber. The strongest weighting by far is given to the fused silica viewports, since their emissivity (estimated at 0.85 ± 0.15) is much greater than the polished steel chamber. Accordingly, 7 out of the 11 sensors are glued directly to the larger fused silica viewports. When we feed the sensor data into our model, we find a BBR shift of 474.6×10^{-17} with an uncertainty of 1.0×10^{-17} for the 3-week ITOC campaign, during which there was a consistent pk-pk inhomogeneity of 800 mK.

The leakage of BBR from the $\sim 600\,\mathrm{K}$ Sr oven is another potential concern. We model this using a similar approach as for the leakage into the cryogenic BBR enclosure in [184], but in our system we find a negligible oven contribution on the order of 2×10^{-19} – the BBR leakage from the oven is strongly suppressed by the presence of two 1 mm-diameter apertures along the atomic beam, before and after the 'spectroscopy cross' described in section 2.2.4. Since I didn't write the BBR models for either the chamber or the oven, I leave a more detailed explanation to their author, William Bowden, to be written up in his doctoral thesis.

We could most likely improve the total temperature inhomogeneity in our chamber to below 200 mK pk-pk if we actively controlled the temperature of the MOT coil cooling water to match the room temperature of around 21 °C – indeed, we will most likely implement active control in the future as part of a new, improved systematic analysis in our current science chamber. However, the performance of our first Sr clock system will probably always be limited by BBR to an uncertainty of at least 2 or 3×10^{-18}. Making further gains will require a new science chamber either with very good thermal design (our preferred option), or with special in-vacuum thermal shields [19] or the silver bullet option, a cryogenic enclosure [184, 271].

Even with perfect temperature control, the inaccuracy of room temperature 300 K Sr lattice clocks is fundamentally limited to a minimum of 1.6×10^{-18} (see table 4.2) by the incomplete data currently available about the Sr atom. To beat this limit will require some significant effort, focused on characterisation of the Sr atom with even more precision. The obvious approach, i.e. taking an improved measurement of the D-manifold lifetime, could potentially reduce the uncertainty in the dynamic shift coefficients. However, even with a perfect lifetime measurement, the 0.2% theoretical uncertainty in the $^3\mathrm{D}_1$ - $^3\mathrm{P}_0$ branching ratio will still limit the dynamic shift uncertainty at around the 1.0×10^{-18} level [230]. Going below 10^{-18} at 300 K will require outside-the-box thinking: perhaps one could undertake a detailed polarisability study of the $^3\mathrm{P}_{0,1,2}$ states (e.g. search for magic wavelengths) in order to narrow down the branching ratio further. Alternatively one might try to map out the BBR shift directly at hotter temperatures up to $\sim 500\,\mathrm{K}$, leveraging the higher sensitivity to dynamic shifts, but the constraints on absolute temperature characterisation during this experiment would be exceptionally tight to achieve a coefficient uncertainty much below 1×10^{-18}.

In all likelihood, the only viable method of reducing the total Sr BBR shift uncertainty to the 10^{-19} level will ultimately be to cool the BBR environment down to cryogenic temperatures.

4.1.2 Zeeman shifts

After BBR, the Zeeman effect introduces the next largest systematic shift in ^{87}Sr lattice clocks. However, its associated uncertainty can usually be kept quite small, as long as the magnetic field is reasonably stable around the Sr atoms.

The first effect we need to worry about is the linear Zeeman shift. For a π-polarised clock probe, this scales as

$$\Delta\nu_B^{(1)} = \frac{\delta g \mu_B M_F B}{h} = \left(\frac{B}{1\,\mu\mathrm{T}}\right) \times 4.890(20)\mathrm{Hz} \tag{4.4}$$

where δg is the differential g-factor between ground and upper state, and the numerical result given is for atoms spin-polarised into the $M_F = 9/2$ state [282]. We observe that the shift coefficient is quite small, because the non-zero δg is only created by the weak $\sim 10^{-4}$ hyperfine-induced admixture of the $^3\mathrm{P}_1$ and $^1\mathrm{P}_1$ manifolds into $^3\mathrm{P}_0$ [44]. As a consequence, it is relatively easy to resolve Hz-level linewidths on the clock transition, regardless of initial M_F state, even without the use of magnetic shielding.

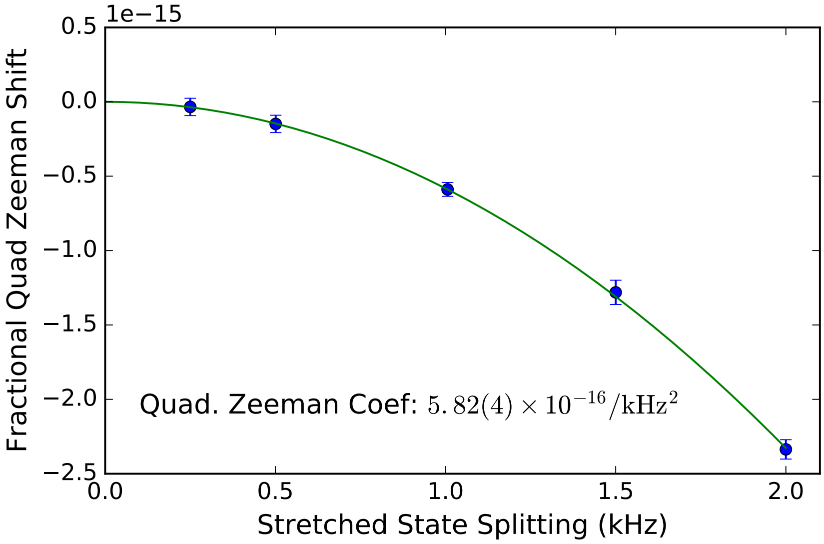

Figure 4.1: Quadratic Zeeman shift data in ^{87}Sr. The y axis has been offset so that the shift curve intersects (0,0) on the graph; in practice, all frequency shifts were all taken with respect to a reference servo operating at 1.006 kHz splitting. The 0.04 uncertainty is purely statistical; systematic effects have not been evaluated.

Nonetheless, the linear Zeeman shift is still quite substantial relative to typical optical clock systematics – if we chose to operate the clock by driving only one M_F sublevel, we'd need to characterise the background magnetic field to within around 1 nT to achieve 1×10^{-17} fractional frequency uncertainty. To circumvent this problem, we always operate the ^{87}Sr clock by combining two clock servos with spin-polarised samples of atoms in the $M_F = +9/2$ and $M_F = -9/2$ states respectively [258]. The clock servo program then constantly calculates an average frequency:

$$\nu_{\text{avg}} = \frac{1}{2}\left(\nu_{+9/2} + \nu_{-9/2}\right) \tag{4.5}$$

for which the linear Zeeman shift completely cancels out. In theory, this cancellation could be compromised by drift in the bias magnetic field between consecutive servo cycles [192]; however, the field stability demonstrated in figure 4.2 implies that our drift-induced residual linear Zeeman shift is at the 3×10^{-19} level or below. If necessary, a re-ordering of servos from $+-+-+-+-...$ to a pattern of $+--+-++-...$ could be implemented to eliminate first-order sensitivity to field drifts.

Since the application of a bias magnetic field is needed in order to resolve the $M_F = \pm 9/2$ sublevels, the quadratic Zeeman shift must also be considered. Both

$M_F = \pm 9/2$ states are shifted in the same direction by the quadratic effect, so the average frequency ν_{avg} changes with bias field according to:

$$\Delta\nu_B^{(2)}/\nu_0 = \left(\frac{B}{100\,\mu\mathrm{T}}\right)^2 \times -5.47(5) \times 10^{-16} = \left(\frac{\nu_{\mathrm{dif}}}{1\,\mathrm{kHz}}\right)^2 \times -5.75(4) \times 10^{-16} \quad (4.6)$$

where the $/\mu\mathrm{T}^2$ coefficient is taken from [284], while the $/\mathrm{kHz}^2$ coefficient is the same weighted average of measurements as used in [193] and is confirmed by the data from our Sr lattice clock shown in figure 4.1[1]. Note that the magnetic field in the right-hand equation is expressed in units of stretched state splitting $\nu_{\mathrm{dif}} = \nu_{+9/2} - \nu_{-9/2}$; we find that the frequency splitting provides a more convenient reference, since it is readily available from the servo data without the need for any kind of magnetic field calibration.

We typically operate our Sr lattice clock with approximately an 894 Hz stretched-state splitting, corresponding to a bias field of 91 µT. This is a compromise between minimising the fractional quadratic Zeeman shift to $46.0(3) \times 10^{-17}$, and maximising the distance between parasitic Zeeman transitions, e.g. $M_F = 7/2 \rightarrow 7/2$, which could potentially introduce line-pulling effects (see section 4.1.8). Our 3×10^{-18} quadratic shift uncertainty is entirely due to the coefficient uncertainty; compared with this, the variations in stretched-state splitting typically give a negligible contribution.

With our ∼ 20 nT magnetic field stability, the Zeeman shift uncertainties will never be a serious problem for the ^{87}Sr clock. An improved measurement of the quadratic shift coefficient could easily be implemented to allow sub-10^{-18} uncertainty at a 91 µT bias field. Meanwhile, improved local oscillator stability will eventually enable the use of smaller bias fields without introducing significant line-pulling effects.

4.1.3 Lattice shifts

We first outlined the concept of a magic-wavelength lattice in section 2.3.3.3: at this wavelength, the polarisabilities of the two clock states are meant to be exactly the same, so the net lattice shift to the clock transition should be zero. However, several residual effects of the lattice light often introduce a significant uncertainty into Sr lattice clocks, comparable to or even larger than the uncertainty of the BBR shift [81, 160, 193, 271].

---

[1] We have not included our own measurement of the quadratic Zeeman coefficient when calculating the clock correction – we only use the published number – because no attempt has been made in our coefficient measurement to characterise potential systematic errors such as might be introduced by lattice vector shifts.

4.1.3.1 Vector and tensor shifts

The hyperfine structure in ^{87}Sr makes the clock transition vulnerable to subtle vector and tensor Stark shifts from the lattice light. These shifts are dependent on the relative orientations of the lattice polarisation $\vec{\epsilon}$, unit k-vector $\vec{e}_k$, and the magnetic field axis $\vec{e}_B$, as well as the internal M_F state of the atoms. Assuming that the bias magnetic field is large enough to set the quantisation axis (typically true), we can write the total scalar, vector, plus tensor Stark shift in the form:

$$\Delta\nu_\alpha = \left(\Delta\nu_s + i\Delta\nu_v\zeta\cos\theta\frac{M_F}{F} + \Delta\nu_t\frac{\left(3\left|\vec{\epsilon}\cdot\vec{e}_B\right|^2-1\right)}{2}\left[\frac{3M_F^2-F(F+1)}{F(2F-1)}\right]\right)\times U_0 \tag{4.7}$$

where $\Delta\nu_s$ is the scalar shift coefficient (note: $\Delta\nu_s = 0$ at the magic wavelength), $\zeta = (\mathbf{e}\times\mathbf{e}^*)$ is the ellipticity of the lattice polarisation, $\cos\theta = \vec{e}_k\cdot\vec{e}_B$, and the vector and tensor shift coefficients near the 813.4 nm magic wavelength are:

$$\Delta\nu_v = 0.22(5)\ \text{Hz}/E_r \qquad\qquad \Delta\nu_t = -4.15(17)\ \text{mHz}/E_r \tag{4.8}$$

where the lattice photon recoil energy $E_r/k_b = 163.3\,\text{nK}$ is used here as a unit of trap depth U_0 [284].

As it turns out, the vector and tensor terms aren't too much of threat to lattice clock accuracy. To begin with, the vector shift is greatly suppressed by the same mechanism as the linear Zeeman effect: since the $M_F = \pm 9/2$ states experience equal and opposite vector shifts, the average frequency $\nu_{\text{avg}} = 1/2(\nu_{+9/2} + \nu_{-9/2})$ remains unchanged. Just as with the Zeeman shift, there might still be room for a systematic error to creep in if the vector shift were to drift significantly between consecutive clock cycles. However, this can be straightforwardly limited in the 1D lattice trap by orientating the quantisation axis perpendicular to the lattice such that $\theta \approx \pi/2$. Then, assuming $U_0 \approx 50E_r$, $|\theta - \pi/2| < 50\,\text{mrad}$, and $\zeta < 0.1$, the vector shift to each stretched state ought to be less than 60 mHz, easily small enough not to worry about residual drift-induced effects.

Unlike the vector shift, the tensor shift is not cancelled by averaging the stretched-state servos. However, it can be efficiently handled using the following strategy:

1. Choose a lattice polarisation where either $\vec{\epsilon}\perp\vec{e}_B$ or $\vec{\epsilon}\parallel\vec{e}_B$, so that $|\vec{\epsilon}\cdot\vec{e}_B|$ is as insensitive as possible to small unintended changes in polarisation.

2. For the chosen polarisation, tune the lattice to an effective magic wavelength $\nu_m + \delta\nu_t$ where the *sum* of the scalar and tensor polarisabilities cancels.

At the effective magic wavelength, the total Stark shift regains its independence of trap depth, making it much more straightforward to limit the lattice shift and its uncertainty. For $|\vec{\epsilon} \cdot \vec{e}_B| = \left[\begin{smallmatrix}0\\1\end{smallmatrix}\right]$, the effective magic wavelength will move by about $\delta\nu_t = \left[\begin{smallmatrix}-100\\+200\end{smallmatrix}\right]$ MHz from the 'true' magic wavelength where the scalar shift $\alpha_s(\nu_m) = 0$ [284]. In fact, this tensor-induced change in the effective magic wavelength may turn out to be an advantage for ^{87}Sr lattice clocks, more easily enabling the construction of multi-dimensional lattice traps based on independent, non-phase-coherent, but fully-magic beams.

The tensor shift could potentially re-emerge as a concern if the lattice polarisation were to drift over time. However, for $U_0 \approx 50E_r$ the tensor shift only varies by around 310 mHz pk-pk over the entire range of polarisations. Therefore, the lattice clock can survive polarisation rotations of up to 80 mrad away from $\vec{\epsilon} \parallel \vec{e}_B$ (or $\vec{\epsilon} \perp \vec{e}_B$) before suffering from a 1×10^{-18} change in tensor shift.

4.1.3.2 Higher-order lattice shifts

So far, we have worked entirely within the electric dipole ($E1$) approximation, and we have only considered 2nd-order perturbations proportional to $|\mathbf{E}|^2$. Although the $|\mathbf{E}|^2$ term is by far the strongest interaction between the lattice light and the atoms, there is also a series of higher-order perturbations which can significantly influence clock accuracy at the 10^{-18} level.

Sticking for now to the $E1$ approximation, we consider the next largest shift: the 4th-order 'hyperpolarisability' proportional to $|\mathbf{E}|^4$. At this order in the electric field, Stark shifts emerge from the off-resonant excitation of two-photon transitions from the clock states, e.g. the two-photon $5s5p\ ^3\mathrm{P}_0$ - $5s4f\ ^3\mathrm{F}_2$ transition at 818.57 nm [48]. For linear polarisation, the hyperpolarisability coefficient has been directly measured in Sr as [284]:

$$\Delta\nu_{\text{hyper}} = \left(\frac{U_0}{E_r}\right)^2 \times (0.46 \pm 0.18)\ \mu\text{Hz} \tag{4.9}$$

Unlike the 2nd-order lattice Stark shift, the hyperpolarisability is strongly dependent on the polarisation ellipticity ζ, regardless of the isotopic nuclear spin. In fact, for atoms like Yb and Hg the hyperpolarisablity even flips sign between linear and circular polarisation, opening up the prospect of a 'magic ellipticity' where the hyperpolarisability is exactly zero [16, 256]. However, Sr has no such sign flip; in fact, the hyperpolarisability is estimated to be roughly 50% larger for circular polarisation than for linear [131], so the coefficient in equation 4.9 is probably as low as it will get. Nonetheless, at typical operating lattice depths the hyperpolarisability is still relatively small in Sr, with $U_0 \approx 50E_r$ leading to a shift of just 3×10^{-18}. At least at the level of our lattice clock error budget in table 4.1, the uncertainty in the hyperpolarisability shift can safely be ignored.

Next, we consider perturbations beyond the electric dipole approximation. The most prominent shifts are due to the off-resonant excitation of magnetic dipole ($M1$) and electric quadrupole ($E2$) transitions, driven respectively by the lattice magnetic field ($|\mathbf{B}|^2$) and electric field gradient ($|\partial_i \mathbf{E}|^2$). At first sight, both of these multipole shifts ought to be linear in the trap depth U_0 – therefore, just as with the tensor shift in section 4.1.3.1, we'd hope they might be easily handled just by moving to a new, slightly adjusted effective magic wavelength.

However, the situation is made more complicated for the standing-wave lattice because of the different spatial dependence of $|\mathbf{B}|^2$ and $|\partial_i \mathbf{E}|^2$ relative to $|\mathbf{E}|^2$ – in fact, in the 1D lattice the magnetic field and electric field gradients have their minimum exactly at the positions where the electric field is at its maximum. This has the important implication that the size of the $M1$ and $E2$ shifts depends on the atomic motional state: hotter atoms will tend to get further away from the electric field maximum than colder atoms, and therefore they will tend to see higher electric field gradients and magnetic fields. Following through on the calculation for an ideal optical lattice and in the harmonic approximation, the $M1/E2$ effects are found to scale as [257]

$$\Delta\nu_{M1/E2} = \xi(n + 1/2) \times \sqrt{U_0} \tag{4.10}$$

where the shift coefficient ξ has been experimentally constrained in [284] to within $\xi = 0 \pm 0.31\,\mathrm{mHz}/\sqrt{E_r}$. At an operating depth of $U_0 \approx 50E_r$, this introduces a fractional systematic uncertainty into the Sr clock of around 3×10^{-18} for ground-state-cooled atoms – small enough to be negligible within our error budget in table 4.1 but definitely large enough to be a concern for state-of-the-art lattice clocks.

The picture gets muddier once we extend the $M1/E2$ model to include hyperpolarisbility and trap anharmonicity. In the full model, we observe a long series of high-order shifts proportional to $U_0^{1/2}$, U_0, $U_0^{3/2}$, and U_0^2, each with different dependences on motional state n, and each estimated to influence accuracy significantly at the 10^{-18} level [131, 207]. These shifts are yet to be experimentally resolved from each other, and so an open question still remains as to how large their coefficients actually

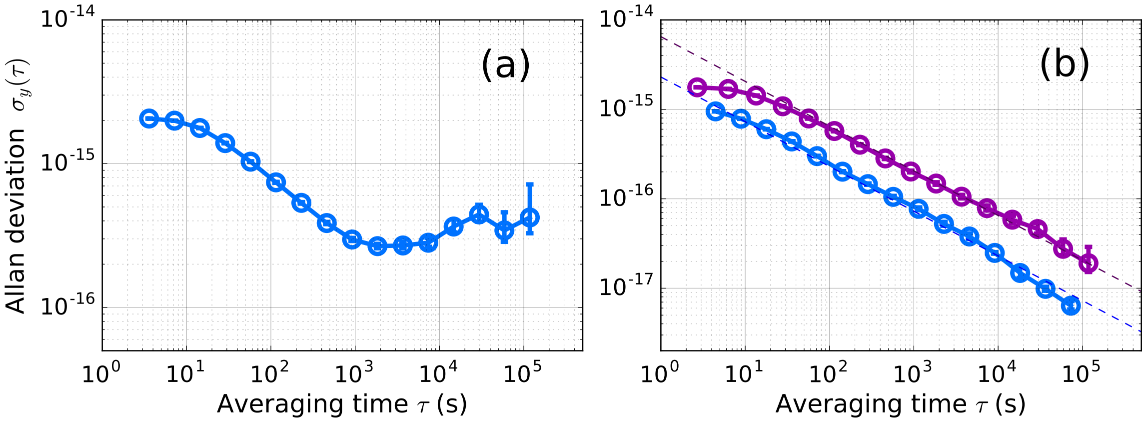

Figure 4.2: (a) The instability in the stretched-state splitting $y_{+9/2} - y_{-9/2}$ for a single continuous dataset. The data implies a long-term background magnetic field stability of around 20 nT. The operating splitting was 997 Hz, or 2.324×10^{-12}. (b) The purple trace shows the instability in the frequency difference $y_B - y_A$ between the two servos of an interleaved self-comparison at different lattice intensities. The blue trace shows the instability of an interleaved self-comparison at different atom densities, with improved stability due to the use of the universal synthesiser to stabilise the local oscillator (see section 3.1.3).

are; however, tackling the proliferation of lattice shifts may well turn out to be the biggest challenge facing optical lattice clocks of all species as they head towards the 10^{-19} range of precision.

4.1.3.3 Total lattice shift evaluation

So far, our Sr lattice clock has only aimed at systematic uncertainties in the 10^{-17} range, which means that we can afford to be relatively rough in our treatment of the awkward non-linear effects. We still, however, need to characterise our linear lattice shift accurately; to do this, we carry out a straightforward extrapolation process:

1. Run an interleaved self-comparison between two servos, operating at different lattice depths U_1 and U_2

2. Calculate the implied lattice shift per unit depth $(\nu_2 - \nu_1)/(U_2 - U_1)$, slightly adjusting ν_1, ν_2 to compensate for the hyperpolarisability shifts at U_1, U_2 as calculated using the coefficient in eqn 4.9

3. Use the result to infer the expected lattice shift at our typical operating depth of $U_0 = 41E_r$ (i.e. where the axial sideband frequency is $\nu_z = 43.6\,\text{kHz}$)

The stability of a representative self-comparison run is shown in figure 4.2; we observe that it takes around a day to reach a statistical measurement uncertainty of 1×10^{-17}. Combining data from a series of self-comparisons up to a maximum depth of $U_{\mathrm{max}} \approx 120E_r$, we infer a null lattice shift of $0(8) \times 10^{-18}$ at $U_0 = 41E_r$, where we have set a lattice polarisation $\vec{\epsilon} \parallel \vec{e}_B$ and a magic lattice frequency 368 554 474(1) MHz.

We take two precautions in order to avoid the contamination of the lattice shift evaluation by collisional effects: (1) we operate the lattice evaluation at low atom number where the total density shift is around 1×10^{-17} fractionally, and (2), following the example of [193], we scale the atom number between the interleaved servos as $U_0^{-3/2}$ in order to cancel out the $\propto NU_0^{3/2}$ dependence of density on lattice depth.

In deeper lattices, care must be also taken to include another effect: atomic motion. In particular, the atoms loaded from the red MOT into deeper lattices tend to be hotter, meaning that they sample less intense regions of the lattice trap. The result, once the lattice shift is plotted against a range of lattice depths, is that we see a 'saturation' effect where the slope of the shift curve levels out as the depth increases. Our method for handling this is always to evaluate the lattice shift against an *effective* depth U_{eff}, given by:

$$U_{\mathrm{eff}} = U_0 - \frac{1}{2}\left(k_B T_z + 2k_B T_r\right) \tag{4.11}$$

where U_0 is the total trap depth inferred from the axial sideband frequency, and T_z and T_r are the axial and radial atom temperatures inferred from the shapes of the sidebands (see section 2.3.3 and [32]). The underlying assumption is that the atoms are approximately in a thermal distribution, and so by the equipartition theorem they ought to posses a potential energy of $\frac{1}{2}k_B T$ along each axis[1]. This assumption is a little speculative, but appears to be quite effective at generating a straight-line dependence of $\Delta\nu_\alpha$ against U_{eff}. Anyway, when the lattice is tuned to the (tensor-adjusted) magic wavelength, the temperature adjustment has little effect – regardless of the correction to U_{eff}, the linear shift per unit lattice depth will always come out as zero at the magic wavelength.

An important limiting feature of our clock is the lack of lattice enhancement cavity around the atoms. We find that this has two consequences. First, it reduces the amount by which we can leverage the lattice intensity, meaning that it takes longer to get good statistics on the Stark shift. Second, it introduces a significant running-wave component of lattice light, meaning that $|\mathbf{B}|^2$ and $|\partial_i \mathbf{E}|^2$ are no longer

---

[1] This U_{eff} method is sourced from private communication with Rodolphe le Targat, SYRTE. It may(!) also be described in his French-language thesis [143].

strictly zero at the **E**-field maxima of the lattice trap. The $|\mathbf{B}|^2$ and $|\partial_i \mathbf{E}|^2$ fields therefore generate $M1/E2$ contributions to the lattice shift which are $\propto U_0$, possibly causing the discrepancy between our measured magic wavelength 368 554 474(1) MHz, PTB's 368 554 465(3) MHz (also not enhanced) [81], and the 368 554 484.9(1) MHz cavity-enhanced magic wavelength measured at JILA [33]. A significant concern for non-enhanced lattices is the possible drift in the retro-reflection alignment, which will have the effect of changing the $M1/E2$ shifts over time. We are yet to evaluate the size of this effect in our Sr lattice clock, but to account for possible drifts (as well as the other non-linear shifts) we inflate the shift uncertainty in table 4.1 from 0.8 to 1.0×10^{-17}.

Looking forward, we will have to investigate the higher-order lattice perturbations much more rigorously in order to assign a shift uncertainty much below 1×10^{-17}. Even in the longer term, evaluation below around 3×10^{-18} in our present system will be a considerable challenge due to the lack of leverage available on the lattice intensity.

In planning a second Sr lattice system aimed at state-of-the-art accuracy, a lattice enhancement cavity will almost certainly be essential: it will allow much greater leverage of lattice intensity, it will exclude the drift-prone running-wave $M1/E2$ contributions, and, with a folding mirror, it will yield better purity of lattice polarisation, fixing down any residual vector/tensor shifts [145]. Thinking more fancifully, one could imagine constructing a cavity-enhanced 3D lattice with beam polarisations specifically oriented for immunity to motion-dependent $M1/E2$ shifts [130], though the required constraints on geometry and polarisation control might turn out to be prohibitive.

4.1.4 DC Stark shift

The DC Stark shift has been characterised very precisely in Sr lattice clocks due to its essential role in determining the BBR shift coefficients (section 4.1.1). The fractional frequency shift can be written [185]:

$$\Delta\nu_{\mathrm{DC}}/\nu_0 = -\frac{1}{2h\nu_0}\Delta\alpha(0)\,|\mathbf{E}|^2 = \left(\frac{|\mathbf{E}|^2}{(1\,\mathrm{V/cm})^2}\right) \times 7.1705(2) \times 10^{-17} \tag{4.12}$$

As discussed in e.g. [174, 192], the experimental DC Stark shift can be measured efficiently by applying external bias fields along each spatial axis. The standard procedure is:

1. Measure the frequency difference $\delta\nu_+$ between two different servos:
   (a) with a positive bias field applied, so that $E_{\text{tot}} = E_{\text{backgnd}} + E_{\text{applied}}$, and
   (b) with no bias field applied, so that $E_{\text{tot}} = E_{\text{backgnd}}$.
2. Measure the frequency difference $\delta\nu_-$ between a different pair of servos:
   (a) with a negative bias field applied, so that $E_{\text{tot}} = E_{\text{backgnd}} - E_{\text{applied}}$, and
   (b) with no bias field applied, so that $E_{\text{tot}} = E_{\text{backgnd}}$.
3. Plug the results into equation 4.12 to infer the DC Stark shift along the working axis. The shift is given by:

$$\Delta\nu_{\text{DC}} = \frac{(\delta\nu_+ - \delta\nu_-)^2}{8\,(\delta\nu_+ + \delta\nu_-)} \tag{4.13}$$

Because the shift is quadratic in $|\mathbf{E}|$, it helps to operate the Sr clock with the background field compensated near $\mathbf{E}_{\text{tot}} \approx 0$ so that small errors $\delta\mathbf{E}$ in the field have a smaller effect. Once the field is approximately compensated, it is perfectly realistic to reach total DC Stark shift uncertainties well below 10^{-18} using equation 4.13 [193].

Unfortunately, the standard approach didn't work too well in our Sr clock. In order to apply the required $\mathbf{E}_{\text{applied}}$ fields, we tried installing external electrode plates as close as possible to the largest viewports on the science chamber. Sadly, even at bias voltages up to ±10 kV we observed no induced shift at the 2×10^{-17} level of frequency uncertainty – it seems that our steel vacuum chamber (section 2.2.4) provides a very effective shield against electric fields!

The failure of the standard approach has left us with a stark choice. We could try applying even bigger voltages to the existing plates. However, there would be no guarantee that we could resolve the shift even at ±30 kV, and at such high voltages we would risk significantly polarising the viewports. Alternatively, we could break vacuum and install internal electrodes much closer to the Sr atoms, but this would be a substantial job – delivery optics would have to be dismantled to gain access for electrode installation, and we would risk compromising the quality of vacuum; re-baking the science chamber would be a huge interruption.

So, instead of trying to salvage the standard approach, we intend to try something a little different: using the atoms themselves as electric field sensors. In order to maximise the atom's sensitivity to electric fields, we plan to excite the atoms into a highly-polarisable Rydberg state as depicted in figure 4.3. The tabulated polarisability data has been given to us by Matthew Jones at Durham University, where

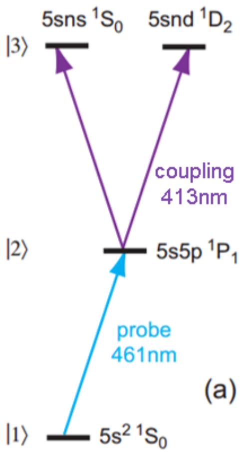

$5snd$	M_J	α (MHz/(V/cm)^2)
56	0	162
56	1	185
56	2	265
75	1	1600
75	2	1700
75	1	2060

Figure 4.3 & Table 4.2: Illustration of the two-photon Rydberg excitation scheme. We intend to measure the frequency of the two-photon transition in order to measure the background DC electric field so that the DC Stark shift can be evaluated.

pioneering research into the Sr Rydberg series has already been under way for a few years [162, 186]. Assuming we can resolve around a 4 MHz spectroscopic line on the $5s^2\ {}^1\mathrm{S}_0$ - $5s56d\ {}^1\mathrm{D}_2$ transition, as achieved in [162], then we ought to be able to use the Stark-induced splitting between the different M_J sublevels to resolve background electric fields as small as $0.2\,\mathrm{V\,cm^{-1}}$, corresponding to a 3×10^{-18} frequency shift to the clock transition.

As yet, the Rydberg measurement has not been taken, and so a question mark remains against the DC Stark shift in our error budget in table 4.1. In all likelihood, given the excellent shielding factor of the steel chamber, the DC Stark shift will turn out to be negligible[1]; anyway, since we have now built an operational 413 nm ECDL, we ought to be finding out soon.

4.1.5 Collisions

As described in e.g. [221], the full many-body physics of collisions in optical lattice clocks is far from trivial. At first sight, the collisional shift in the spin-polarised, fermionic ^{87}Sr clock ought to be very small indeed – s-wave collisions should be suppressed by the Pauli exclusion principle, whereby indistinguishable fermions must

---

[1] One caveat is the presence of a vestigial BBR tube described in [106], which protrudes into the chamber with a copper/graphite tip around 3 cm away from the lattice-trapped atoms. The entire tube is nominally conductive, but any oxidisation of the copper may cause patch charges which could induce an electric field at the atoms.

be antisymmetric under pairwise exchange, and so only odd orders of the partial wave expansion (p-wave, f-wave, and so on [280]) should be allowed. Meanwhile, any collisions of a higher order than s-wave ought to be effectively frozen out, with the $\sim 2\,\mu\mathrm{K}$ atoms being much too cold to breach the estimated $\sim 75\,\mu\mathrm{K}$ centrifugal barrier for p-wave collisions [175].

Nonetheless, substantial collisional shifts have been discovered in fermionic lattice clocks with both $^{87}\mathrm{Sr}$ [167] and $^{171}\mathrm{Yb}$ [152], prompting a string of investigations into their cause. One culprit is inhomogeneity in atomic excitation, which undermines the indistinguishability of the atoms and therefore lets a strong s-wave collision channel emerge [50, 86]. Such inhomogeneity can be caused by probe intensity gradients, but the dominant cause is often the spread in Lamb-Dicke parameters between atoms in different motional states (see section 2.3.3) – the effect can therefore be minimised by reducing the atom temperature and by ensuring that the probe is well aligned along the lattice axis [32].

Beyond the residual s-wave interaction, p-wave effects must also be taken into account in the most accurate lattice clocks. Elastic p-wave collisions have been directly observed in fermionic $^{171}\mathrm{Yb}$ [151], while full theoretical models including both s-wave and p-wave interactions in the $^{87}\mathrm{Sr}$ clock [221] have been corroborated by experimental observations of frequency shifts, decoherence, and inelastic losses [29, 175, 176].

Regardless of the detailed physics underlying them, collisions present two major problems for the clock-maker: (1) they can cause decoherence on the clock transition, limiting the viable length of clock interrogation pulses at high atom numbers [175], and (2) they introduce a density-dependent shift to the clock transition.

In our lattice clock, the comparatively short 80 ms Rabi pulse time means that we can ignore the decoherence effects. However, we still need to address the collisional shift. As a first step, we simply operate at the lowest density we can – since our clock stability is limited by the Dick effect at around $\sigma_y(\tau) \approx 2 \times 10^{-15}/\sqrt{\tau}$ (see chapter 3 and figure 4.2), the quantum projection noise remains irrelevant even with as few as $N \approx 10^2$ atoms. In addition, we keep atom density to a minimum by maximising the volume of the sample – the fairly wide $65\,\mu\mathrm{m}$ lattice waist spreads the atoms out radially, the low lattice $41E_r$ trap depth spreads them out in all three dimensions on each site, and a diffuse red MOT (low field gradient and large detuning) loads as many lattice sites as possible.

After these shift-minimisation precautions, we characterise the density shift by linear extrapolation. The process is almost identical to the lattice shift evaluation:

1. Run an interleaved self-comparison between two servos, operating at different atom numbers N_1 and N_2
2. Calculate the implied density shift per unit atom number $(\nu_2 - \nu_1)/(N_2 - N_1)$
3. Use the result to infer the expected density shift at a typical operating atom number N_0

Using this extrapolation technique, we evaluate a density shift of $8(8) \times 10^{-18}$ in fractional units. The uncertainty on this evaluation is minimised by modulating the atom number with a large lever arm of $N_2 - N_1 \approx 5N_0$ or more; nonetheless, no shift has yet been resolved at typical N_0[1]. We are yet to calibrate the absolute atom number N_0 – the relative numbers, measured using the fluorescence detection technique described in section 2.2.4.2, are sufficient for inferring the density shift. We note that an implicit assumption in the extrapolation is that the shift is linear in atom number; this isn't foolproof, but it's plausible in the weakly-interacting regime [192], especially to within the 100% uncertainty reached in our evaluation.

At the 1×10^{-17} uncertainty presently accessible to us, the density shift is clearly not a limitation. However, looking further ahead, density effects may become significantly more problematic. For example, in a Dick-free lattice clock of the type described in section 3.4, we would expect the clock stability to be limited by quantum projection noise rather than the Dick effect. In this regime, low 10^2 atom numbers would sharply limit the clock performance – in such a system, increasing the atom number to the experimentally-plausible 10^6 range would let us gain a factor of 100 in measurement precision. Meanwhile, collisional decoherence has been found to undermine stability even in the present generation of clocks, limiting the maximum viable length of clock interrogation pulse [175]. With these considerations in mind, a more complete solution to suppress collisional effects in lattice clocks would be worth pursuing.

A brute-force method of reducing collisions would be simply to increase the trap volume further. This method was successfully implemented in [194] using an enormous 160 µm cavity-enhanced lattice waist, and could be taken even further by also exploring more diffuse loading techniques such as the blue MOT 'atom drain' in [48, 144]. However, these increases in trap volume are only likely to suppress the collisional effects by one or two orders of magnitude at most. If we want a viable route towards $\sim 10^{-19}$ accuracy with $\sim 10^6$ atoms, we will have to look elsewhere.

---

[1] Earlier density shift measurements did resolve a shift of a few parts in 10^{16}, based on an order of magnitude more atoms, a deeper lattice $U_0 \approx 70E_r$, and a tighter 45 µm waist.

Paradoxically, the collisional effects might be most efficiently suppressed by using a *higher* atomic density. This approach was pursued in [28], where dense samples of ^{87}Sr were trapped in a tightly-focused 2D lattice, entering a strongly interacting regime where the collisions were manifested as resolvable sidebands. The strongly interacting regime could be exploited by locking the clock only to the carrier transition, ultimately realising an accurate ^{87}Sr lattice clock with reduced sensitivity to density-dependent shifts and decoherence [251, 252]. Going one step further, the ultimate solution may be to use a large-volume, cavity-enhanced 3D lattice. Here, the collisional sidebands in doubly-occupied sites should be separated further out from the carrier than was seen in the 2D lattice. In addition, a maximum occupancy of one atom per site could be imposed by implementing a photoassociation spectroscopy pulse on the narrow-line cooling transition at 689 nm prior to clock interrogation, so that any pairs of atoms in the same site will form molecules and fall out of the trap [4].

4.1.6 Doppler shifts

Care must be taken to avoid Doppler shifts in all species of atomic clock: even an atomic velocity as small as $0.3\,\mathrm{nm\,s^{-1}}$ is enough to induce a 10^{-18} Doppler shift on the clock transition. In optical lattice clocks with sufficiently deep lattices, each atom's position will accurately track a single antinode of the lattice trap [153], implying that the Doppler shift will be primarily determined by the amount of relative shaking (or, equivalently, optical phase chirping) between the clock probe and the lattice beams.

As a first step towards avoiding Doppler shifts, we passively minimise all sources of mechanical vibration on the optical table: all laser beam shutters use very lightweight piezo-actuated flags, while the bulkier camera shutter (Uniblitz VS14S2T0-NL) – the only audible object switching during the sequence – is mounted on thick, vibration-damping rubber feet. The residual vibrations introduced e.g. by the switching MOT coils are further suppressed by the rigid construction of all beam delivery optics surrounding the science chamber (we use plenty of 1.5″ steel posts).

In addition to the passive avoidance of vibrations, we also actively phase-stabilise the clock probe light delivered to the atoms [170] (see also section 3.1.1 for an outline of the overall clock laser setup). The clock phase is stabilised at a reference sampler placed as close as possible to the lattice retro-reflecting mirror, cancelling out any fluctuations in the optical path up to that sampler. Meanwhile, vibrations in the optical path *after* the lattice retro-reflector should affect the phases of the clock probe and the lattice light by (almost) the same amount, enforcing a zero net Doppler shift. Therefore, vibration-induced Doppler shifts can only emerge within the short 20 cm uncompensated path between the reference sampler and the lattice retro-reflector – in theory, the phase noise from such a short path ought to be negligible [80].

Apart from vibrations, the final acousto-optic modulator (AOM) used for switching the clock probe beam is another infamous source of potential phase chirps – and thus Doppler shifts – in optical clocks [65, 80, 175]. The main problem arises from the RF-induced thermal heating of the AOM crystal: the heating causes the crystal to expand dynamically during the clock pulse, injecting a phase chirp onto the first-order-diffracted beam used for atomic interrogation. (Additional transient phase excursions also occur at the beginning and end of the pulse due to the bandpass RF-to-optical transfer function [65], but these should settle within around 100 ns in our TeO_2 AOM). In order to suppress the thermal AOM phase chirp, we follow the same approach as in [175]: the sampler for path length stabilisation is placed after the AOM, and reflects the 0^{th} order beam back onto itself to generate the phase-stabilising signal. Technically, this has the effect of phase-stabilising the 0^{th} order beam, not the 1^{st} order beam actually used for atomic interrogation; however, with a deflection angle of just 30 mrad, both beams ought to sample the same path length within the AOM within less than 0.1%. The effectiveness of the technique has been borne out by measurements in [175], but we are yet to measure the chirp suppression achieved in our setup – we will, soon!

We quote a rather pessimistic uncertainty of 2.9×10^{-17} for the Doppler shift in table 4.1. This is based on an an interleaved self-comparison between the 'normal' clock servo, and a test servo operated with an extra ~ 100 ms delay inserted between the end of the red MOT and the beginning of the clock probe pulse. The idea behind this self-comparison is that if vibrations are induced by the MOT coils switching or by anything in the cooling sequence, they ought to have settled (or evolved in phase) within the extra ~ 100 ms delay time. As expected, a null shift was measured, but the statistical uncertainty in this measurement was 2.9×10^{-17}, which we cautiously quote as an uncertainty in the systematic shift. In order to refine this uncertainty, we need to: (1) directly measure the residual phase chirp on the final switching AOM using a beat note against an independently-delivered beam, and (2) take a longer dataset to verify the null shift to a greater precision. Once these steps are implemented, we do not foresee the Doppler shift as a serious impediment to 1×10^{-18} clock accuracy.

As a side note, we observe that the thermal velocity of the Sr atoms at 2 μK is only around $0.02\,\mathrm{m\,s^{-1}}$, implying a second-order Doppler shift well below 10^{-20}; this contribution is decidedly negligible within any foreseeable uncertainty budget.

4.1.7 Probe Stark shift

We only need a tiny probe intensity at 698 nm to drive an 80 ms Rabi pulse on the ^{87}Sr clock transition, but the Stark shift from this light still isn't completely negligible. Here, we calculate the shift theoretically, using the estimated hyperfine-induced natural linewidth of 1.2 mHz [213, 232], our calculated differential polarisability of $\Delta\alpha(698.4\,\mathrm{nm}) = 607.0$ atomic units, and the Clebsch-Gordan coefficient of 0.9045 for the $M_F = 9/2 \rightarrow 9/2$ transition. From these numbers, we infer a probe Stark shift of:

$$\Delta\nu_{\mathrm{probe}} = \left(\frac{T_p}{100\,\mathrm{ms}}\right)^2 \times 1.3 \times 10^{-18} \tag{4.14}$$

where T_p is the Rabi π-pulse probe time.

However, equation 4.14 only holds for perfectly motionless atoms. In reality, a lattice depth of $U_0 = 41E_r$ yields a finite Lamb-Dicke parameter of around $\eta_z \approx 0.33$ (see section 2.3.3.2), and if we also account for imperfect probe alignment and trap anharmonicity, we might expect to need up to 1.5× more probe power to drive the same Rabi frequency. For our 80 ms Rabi π pulse, this yields a probe Stark shift of 3×10^{-18}, which we quote only as an uncertainty in our error budget.

Since the probe Stark shift can be efficiently reduced by going to longer interrogation times T_p, we anticipate that this contribution to the error budget will easily be manageable below 1×10^{-18} in the future.

4.1.8 Line-pulling

Line-pulling can occur from accidental, off-resonant excitation of a transition other than $M_F = 9/2 \rightarrow 9/2$ carrier. If the accidental sideband excitation is imbalanced between the two interrogation points of the Rabi locking scheme at $\nu_0 \pm \Delta\nu_{\mathrm{FWHM}}/2$, then it adds an offset to the error signal, therefore nudging the locked clock into the wrong equilibrium frequency.

We first consider the possible line-pulling from motional sidebands. Since the axial lattice sidebands are separated from the carrier by more than 40 kHz, much greater than the transition linewidth of $\Delta\nu_{\mathrm{FWHM}} = 10\,\mathrm{Hz}$, we can safely ignore their line-pulling influence. The first-order radial lattice sidebands are expected to be much closer to the carrier, at approximately ±120 Hz, but: (1) they are so small we have never been able to see them, and (2) the high radial quantum number $\bar{n}_r \sim 350$ means that the Rabi frequencies of the nearest red/blue radial sidebands should be balanced at roughly the $\sqrt{351}/\sqrt{350} \approx 1.001$ level (see equations 2.27). A final, more subtle effect of atomic motion is tunnelling, which theoretically generates sidebands at around ±870 Hz due to the gravitational potential energy difference between neighbouring vertical lattice sites; however, even accounting for a reduced site-to-site barrier due to radial atomic motion, we estimate the tunnelling shift to be well below 10^{-19} for our operating depth of $U_0 = 41E_r$ [153][1]. We can, therefore, safely rule out line-pulling from all types of motional sidebands.

The bigger threat is from the accidental excitation of undesired Zeeman transitions. We expect the dominant contribution to come from the closest sideband, $M_F = 7/2 \to 7/2$, arising from imperfect spin-polarisation of the ^{87}Sr sample. We model the accidental excitation of this sideband by assuming that it undulates beneath a Lorentzian envelope:

$$F_{\text{env}}^{(7/2)} = \frac{C}{1 + \left(2(\nu_{\text{probe}} - \nu_{7/2})/\Delta\nu_{\text{FWHM}}\right)^2} \tag{4.15}$$

where C, $\nu_{7/2}$, and $\Delta\nu_{\text{FWHM}}$ are respectively the contrast, the frequency, and the linewidth of the $M_F = 7/2 \to 7/2$ transition.

To estimate the worst-case line-pulling, we assume that sideband excitation has a maximum at the 'low'-side interrogation frequency (i.e. $F_{\text{low}}^{(7/2)} = F_{\text{env}}^{(7/2)}$), and has a zero on the 'high' side (i.e. $F_{\text{high}}^{(7/2)} = 0$). These assumptions put an upper limit on the accidental excitation imbalance of:

$$\frac{1}{2}\left(F_{\text{high}}^{(7/2)} - F_{\text{low}}^{(7/2)}\right) < \frac{1}{2}\left(\frac{0.1}{1 + (2 \times 84/10)^2}\right) = 1.8 \times 10^{-4} \tag{4.16}$$

where we have inserted our operating Zeeman splitting of 89 Hz (minus the hopping frequency $\Delta\nu_{\text{FWHM}}/2 = 5\,\text{Hz}$), and we have pessimistically assumed a contrast of $C = 0.1$ on the $M_F = 7/2 \to 7/2$ transition, even though the spin-polarisation efficiency is typically 95% or more.

Using the equations derived in section 3.2, we can convert the excitation imbalance in equation 4.16 directly into a frequency offset for the Rabi lock: the frequency discriminant is given by $dF/d\nu = 1.9T_p = 0.152\,\text{Hz}^{-1}$ for our Rabi pulse time $T_p = 80\,\text{ms}$, yielding a fractional frequency shift of:

$$\Delta\nu^{(7/2)}/\nu_0 = \frac{1}{\nu_0}\left[\frac{\frac{1}{2}\left(F_{\text{high}}^{(7/2)} - F_{\text{low}}^{(7/2)}\right)}{dF/d\nu}\right] < 3 \times 10^{-18} \tag{4.17}$$

---

[1] The tunnelling shift scales exponentially with lattice depth, and quadratically with pulse time. Longer 1 s pulse times may safely enable the use of shallower $\sim 20E_r$ lattices, thereby reducing lattice Stark shifts; however, a proper experimental investigation into tunnelling might yet be worthwhile.

We note that the frequency discriminant for the Rabi lock scales as T_p, while the sideband linewidth scales as T_p^{-2}, implying that the line-pulling shift strongly scales with pulse time as T_p^{-3}. Line-pulling from some other sidebands, such as $M_F = 7/2 \rightarrow 9/2$, scales more weakly as T_p^{-2} because their accidental excitation can coherently interfere with the carrier excitation; however, in our clock we estimate that these coherent sidebands don't cause significant offsets – for example the $\sigma^{\pm}$ Zeeman transitions are suppressed by the approximate π polarisation of the clock probe light, as well as by their larger detuning from the carrier.

Despite our reassuring theoretical results for the line-pulling shift, we aren't completely confident in disregarding it as a problem for the ^{87}Sr lattice clock. Our main concern is that the theoretical results assume that the probe light is exactly monochromatic; however, in reality the local oscillator possesses an unknown spectrum of frequency noise sidebands – this noise could potentially excite the Zeeman sidebands much more than a strictly monochromatic probe would.

As yet, we haven't been able reliably to characterise the laser-noise-induced line-pulling. As a preliminary investigation, we ran an interleaved self-comparison between one servo with a bias magnetic field of $\nu_{9/2} - \nu_{-9/2} = 997\,\text{Hz}$ and another at 1101 Hz. This comparison was intended as a 'shot in the dark' for line-pulling effects – we theorised that if there is any noise-induced line-pulling from the $M_F = 7/2 \rightarrow 7/2$ transition, then that line-pulling shift ought to change along with the Zeeman splitting. As hoped, the comparison yielded a null shift when adjusted for the quadratic Zeeman effect. For lack of a more convincing upper bound, the 2.8×10^{-17} statistical uncertainty on this measurement is quoted as an uncertainty in the line-pulling shift in the error budget table 4.1. In all likelihood, the line-pulling shift is actually much smaller.

Going forward, there are a few possible options for further reducing the line-pulling uncertainty:

- Implement an extra state-preparation stage prior to the main Rabi pulse, consisting of a short pulse on the $M_F = 9/2 \rightarrow 9/2$ clock transition followed by a flash at 461 nm to blast away any atoms remaining in the ground state [81, 282]. In combination with a systematic optimisation of probe π polarisation, this could all but eliminate the Zeeman sidebands.

- Precisely measure the quadratic Zeeman shift coefficient (we can apply equally stable fields up to $\nu_{9/2} - \nu_{-9/2} \approx 30\,\text{kHz}$ for this purpose), and then apply a larger bias field during ordinary clock operation in order to distance the Zeeman sidebands further from the carrier.

- Use a longer Rabi pulse time T_p to suppress sideband excitation.

Frustratingly, none of these approaches are strictly rigourous. It would be preferable to measure the laser noise spectrum directly, but I can't think of a way to do this with sufficient precision.

4.1.9 Servo error

The clock servo tries its best to make sure that the local oscillator is constantly steered exactly onto atomic resonance, but no servo is ever perfect. Typically, lock offsets can emerge from two main sources: (1) finite servo gain, which can allow the local oscillator to 'sag' above or below the clock transition frequency [209], and (2) asymmetric sampling of the local oscillator drift, which is left uncompensated in between AOM frequency updates.

As mentioned in section 3.2.1, our clock servo uses two stages of integration to minimise servo sagging. To recap, we correct the laser frequency with the sum of an integrated component $\propto \sum\limits_n \epsilon_n$ and a doubly-integrated component $\propto \sum\limits_m \left(\sum\limits_{n=1}^{m} \epsilon_n \right)$. The role of the higher-gain single integrator is to capture lock stably to the atomic transition within a few clock cycles ($\sim 10\,\mathrm{s}$), while the slower second integrator is intended to compensate the local oscillator's linear drift in the longer run. Once the both integrators have successfully settled, which takes roughly 100 s, the leading contribution to residual lock offset is the small non-linear drift in the local oscillator frequency.

We analyse the residual servo error in post-processing. Whenever running the clock, we save the measured error signals ϵ_n in each clock cycle to a text file (along with other data such as atom number, AOM frequencies etc.). To estimate the servo error during a series of N cycles, we convert the time-averaged error signal ϵ_n into an implied average frequency offset using the frequency discriminant $d\epsilon_n/d\nu_n$:

$$\Delta\nu_{\mathrm{servo}} = \frac{1}{N} \sum_{n=1}^{N} \frac{\epsilon_n}{d\epsilon_n/d\nu_n} \tag{4.18}$$

where $d\epsilon_n/d\nu_n = 1.9T_p = 0.152\,\mathrm{Hz}^{-1}$ for our Rabi lock with $T_p = 80\,\mathrm{ms}$. Our 2×10^{18} servo error quoted in the table 4.1 is calculated using equation 4.18 for a worse-than-average run of data. On a good day, the servo error is significantly smaller.

An implicit assumption underlying equation 4.18 is that the average error signal is an accurate reconstruction of the average local oscillator frequency error, but this assumption is only true if the local oscillator remains within the linear region on the sides of the Rabi fringe. This caveat is one important reason that we use a fairly conservative pulse length of 80 ms, despite our capability to lock reliably to Rabi features up to $T_p \approx 200\,\mathrm{ms}$ – if the local oscillator were to have regular excursions on the order the transition linewidth $\Delta\nu_{\mathrm{FWHM}}$ (as it might for the longer Rabi pulses), then the error signal would become deeply non-linear in frequency, and therefore could not be used to reconstruct the average local oscillator frequency error.

Finally, we note that in order to avoid asymmetric sampling of the local oscillator drift (roughly $100\,\mathrm{mHz\,s^{-1}}$ in our cavity), we consistently apply the 2nd integrator corrections at short 5 ms intervals – the linear cavity drift is therefore only very briefly left uncompensated, with frequency excursions typically limited to a negligible 0.5 mHz level.

4.1.10 Background gas collisions

We follow [87] to approximate the frequency shift from background gas collisions. We assume that the background gas is mostly hydrogen, so we can use the Sr-H C_6 coefficients calculated in [188] – the validity of this assumption is aided by the similar C_6 coefficients for the other likely background gases. Using our measured $1/e$ vacuum lifetime of $1/\Gamma_{\mathrm{vac}} = 8\,\mathrm{s}$, we calculate a background shift of

$$\frac{\Delta\nu_{\mathrm{backgnd}}}{\nu_0} \approx \frac{1}{\nu_0}\frac{\Gamma_{\mathrm{vac}}}{\pi}\frac{\Delta C_6}{13.8 C_6} \approx 2\times10^{-18} \tag{4.19}$$

which is in good agreement with the estimate in [81] for a similar vacuum lifetime. We list the shift from background gas collisions as an uncertainty in table 4.1.

More care will need to be taken to characterise background gas collisions as lattice clocks approach 1×10^{-18} total uncertainty. In particular, the substantial discrepancy ought to be resolved between the 7×10^{-20} shift estimate for the SrI system in [33] and the 1.5×10^{-18} estimate for the Sr system in [81], despite their very similar reported vacuum lifetimes of 8 s [175] and 7 s respectively. The resolution of this discrepancy may have important implications in inflating the rather low 6×10^{-19} shift uncertainty quoted for the world-leading SrII system in [193], in which the vacuum lifetime was reportedly only $\sim 1\,\mathrm{s}$ [174].

4.2 Prospects for the ^{88}Sr lattice clock

The recent progress of optical lattice clocks towards 10^{-18} accuracy and stability has been based primarily on *fermionic* isotopes of the alkaline-earth metals, such as ^{87}Sr [81, 145, 193, 271], ^{171}Yb [109], and ^{199}Hg [178, 294]. Alongside the fermionic clocks, there has also been a lesser effort to realise lattice clocks with bosonic isotopes, e.g. ^{88}Sr [5], ^{174}Yb [211], and ^{24}Mg [141]. So far, however, the performance of these bosonic clocks has lagged behind, with the best reported accuracies only at the 10^{-15} level.

As we initially outlined in section 2.1.1, there are two main reasons that fermions ($I \neq 0$) have often been preferred over bosons ($I = 0$):

1. The $I \neq 0$ hyperfine interaction in the fermion admixes a small fraction of ^{1}P$_1$ and ^{3}P$_1$ into ^{3}P$_0$, introducing a significant natural linewidth $\gtrsim 1\,$mHz to the clock transition, and thus conveniently allowing direct excitation with minimal probe Stark shifts.

2. The Pauli exclusion principle for fermions suppresses the low-energy s-wave collision channel, meaning that collisional decoherence and frequency shifts can be 'frozen out' by laser cooling the atomic sample to µK temperatures.

These features put the fermionic clocks at an advantage, but there are still a handful of reasons to revisit neglected bosons like ^{88}Sr.

Firstly, the bosonic isotopes are often much easier to capture, cool, and state-prepare because of their simpler atomic structure. An extreme example is the Mg lattice clock, where the fermionic ^{25}Mg isotope suffers from strong leakage into dark hyperfine states in the 2$^{\text{nd}}$-stage MOT, leaving bosonic ^{24}Mg as the only isotope yet successfully trapped in a lattice [141]. The case of Sr is less extreme, but the boson is still significantly easier: compared with ^{88}Sr, we need an extra 'stirring' laser to operate the ^{87}Sr red MOT, and we need significantly more precise frequency control at 689 nm to load a repeatable number of ^{87}Sr atoms into the lattice (see section 2.3.2). In addition, ^{88}Sr has an eleven times greater natural abundance, is easier to repump, and needs no spin-polarisation stage before clock interrogation. All these factors mean that ^{88}Sr easier to work with, making it a good starting point for newly-developed or transportable systems [34, 204]. Even in systems capable of loading ^{87}Sr, the fact that we can quickly load around 100 times as many ^{88}Sr atoms into the lattice is a good reason to choose the boson: it enables the reduction of dead time, and thus also the Dick-effect instability (see section 3.3.4), and it allows us to use much shallower lattice traps (with correspondingly lower lattice shifts) while still retaining enough atoms to operate the clock.

Another advantage of bosons is the avoidance of certain systematic shifts. In particular, the vector and tensor lattice shifts are no longer present, avoiding the need for careful lattice polarisation control (see section 4.1.3.1); in addition, there is no danger in the boson of line-pulling from neighbouring Zeeman sublevels (see section 4.1.8), since there is genuinely only one $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_0$ clock transition. Finally, the boson has no first-order Zeeman shift – thus, *if* no bias magnetic field were required, the bosonic lattice clock would see negligible sub-10^{-18} total Zeeman shifts even with background field fluctuations at the several µT level.

Partly for these reasons, and partly just to explore something a little different from the other lattice clock groups, we devote the rest of this section to a discussion of the forlorn ^{88}Sr isotope. We start in section 4.2.1 by introducing the main challenges of building an accurate ^{88}Sr clock using the conventional technique of magnetically-induced spectroscopy. Then, in section 4.2.2, we explain the theory behind our 'modified hyper-Ramsey' scheme for the elimination of probe Stark shifts, including an experimental demonstration of its effectiveness on ^{88}Sr. Finally, in section 4.2.3, we outline some multi-photon excitation schemes for bosonic lattice clocks which could enable them to fulfil their potential as truly magnetically-insensitive clocks.

4.2.1 Challenges for the ^{88}Sr clock

The main challenges facing ^{88}Sr lattice clocks are down to the absence of the more convenient features of ^{87}Sr:

1. Since the $J = 0 \nrightarrow J = 0$ selection rule remains unbroken by nuclear spin, the $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_0$ clock transition in bosonic lattice clocks is completely forbidden to all orders in the multipole expansion. While this makes the natural linewidth exceptionally narrow ($1/\gamma_{\mathrm{nat}} \approx 5800$ years) [232], it also makes it impossible to excite the clock transition with a single photon.

2. Since *s*-wave collisions are allowed by the bosonic spin statistics, the clock transition suffers from significant density-dependent decoherence and frequency shifts, even at ultracold atom temperatures [161].

In the face of the strict $J = 0 \nrightarrow J = 0$ selection rule, the only tried and tested method exciting the clock transition in bosonic lattice clocks is through magnetically-induced spectroscopy [5, 15, 254]. According to this technique, a magnetic 'mixing'

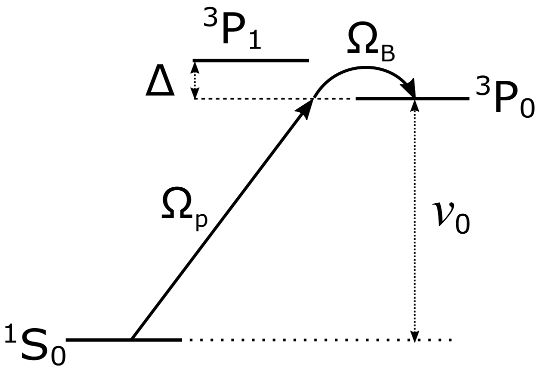

Figure 4.4: The basic principles of magnetically-induced spectroscopy (see text).

field provides a zero-frequency $M1$ photon so that the $^1\mathrm{S}_0 \xrightarrow{E1} {}^3\mathrm{P}_1 \xrightarrow{M1} {}^3\mathrm{P}_0$ Raman transition can be driven with a single probe laser (i.e. the 698 nm clock laser for ^{88}Sr).

This excitation mechanism is well described by first-order perturbation theory. The magnetic mixing field introduces a Hamiltonian $\hat{\mu} \cdot \mathbf{B}$ to the atoms, where $\hat{\mu} = g_L \mu_B \hat{\mathbf{L}} + g_S \mu_B \hat{\mathbf{S}}$, which induces a first-order perturbation to the upper clock state [17]:

$$\left|{}^3P_0'\right\rangle = \left|{}^3P_0\right\rangle + \frac{\Omega_B}{\Delta}\left|{}^3P_1, M_J = 0\right\rangle \tag{4.20}$$

where $\Omega_B = \left\langle {}^3P_1, M_J = 0\right| \hat{\mu} \cdot \mathbf{B} \left|{}^3P_0\right\rangle \approx \sqrt{2/3}\mu_B|B|/\hbar$, and $\Delta/2\pi = 5.6\,\mathrm{THz}$ is the frequency splitting between ^{3}P$_0$ and ^{3}P$_1$.

If the ^{88}Sr atoms are then subjected to a resonant probe beam $\mathbf{E}_p \cos \omega_0 t$ with polarisation $\mathbf{E}_p \parallel \mathbf{B}$, we find that the overall clock transition is driven with a Rabi frequency:

$$\Omega_{\text{clock}}/2\pi = \frac{\Omega_B \Omega_p}{2\pi\Delta} = \alpha|\mathbf{B}|\sqrt{I_p} \tag{4.21}$$

where $\Omega_p = \left\langle {}^1S_0\right| \hat{\mathbf{d}} \cdot \mathbf{E}_p \left|{}^3P_1, M_J = 0\right\rangle/\hbar$, I_p is the probe laser intensity, and the various constants of proportionality have been gathered into the single factor α, which in Sr follows $\alpha = 6.26\,\mathrm{Hz}/(\mathrm{mT}\sqrt{\mathrm{W\,cm^{-2}}})$ [254].

Unfortunately, it turns out that we need to apply somewhat large magnetic fields $|\mathbf{B}|$ and probe intensities I_p in order to drive the clock transition within a reasonable probe time. These strong fields introduce substantial probe Stark and quadratic Zeeman shifts to the clock transition according to:

$$\Delta\nu_B^{(2)}/\nu_0 = \left(\frac{B}{1\,\mathrm{mT}}\right)^2 \times -5.47(5)\times 10^{-14} \tag{4.22}$$

$$\Delta\nu_p/\nu_0 = \left(\frac{I_p}{1\,\mathrm{W/cm^2}}\right) \times -6.6\times 10^{-14} \tag{4.23}$$

where we use the Zeeman shift coefficient measured in ^{87}Sr [282] which should be the same as the ^{88}Sr coefficient well within the given error, and we use our own calculated Stark shift coefficient according to the methods described in [106][1].

In order to get a sense of the magnitude of these probe-related shifts, we can combine equations 4.21, 4.22 and 4.23 into the form:

$$\frac{\sqrt{\Delta\nu_B^{(2)}\Delta\nu_p}}{\nu_0} = \left(\frac{100\,\mathrm{ms}}{T_p}\right) \times 4.8\times 10^{-14} \tag{4.24}$$

where $T_p = \pi/\Omega_{\mathrm{clock}}$ is the time it takes to drive a Rabi π pulse on the clock transition.

Equation 4.24 implies that, if we choose roughly to balance the Zeeman and Stark effects so that $\Delta\nu_B^{(2)} \approx \Delta\nu_p$, then both shifts will be at the 10^{-14} level for realistic Rabi pulse times. Controlling such substantial shifts is a huge challenge! In addition, we observe that some of the nominal advantages of ^{88}Sr are strongly compromised by operating with such a large mixing field – the sensitivity to background magnetic field fluctuations is worsened, even relative to ^{87}Sr, and vector lattice Stark shifts are subtly reintroduced through third-order $M1E1E1$ perturbations [253].

The main topic of the following sections is the mitigation of probe-related shifts in ^{88}Sr: we outline how 'hyper-Ramsey' methods can be used to suppress the probe Stark shift in section 4.2.2, and in section 4.2.3 we propose alternative multi-photon excitation routes which promise totally to avoid the Zeeman shift. As we shall find, these methods could facilitate control of the total ^{88}Sr probe shift at the 10^{-18} level.

Before addressing the probe-related shifts, we briefly turn to the problem of collisions in ^{88}Sr, basing our discussion on the coefficients measured in [161]. Plausibly, the two approaches discussed in section 4.1.5 are as applicable to minimising the collisional shift in ^{88}Sr as they were in ^{87}Sr. In the first approach, we simply reduce the atomic density as far as possible in the standard, 1D lattice trap. Plausibly, we could streak load the ^{88}Sr atoms along a 1 mm length of lattice sites (see section 2.3.2), which combined with a 150 µm waist and a shallow lattice depth $U_0 \approx 30E_r$ would

---

[1] The discrepancy with the Stark shift coefficient given in [254] is due to the approximations made in that paper – our value is more accurate, and in good agreement with [215].

yield a manageable collisional shift of around 3×10^{-18} per 1000 atoms [161]. Alternatively, a high-density approach may be most effective, especially at higher atom number: following the example of [4, 5], one could load into a 3D lattice and clean up the doubly-occupied sites using photoassocative pulse on the red MOT cooling line. If we then probe such a system with a sufficiently stable local oscillator, collisions in residual doubly-occupied sites might be resolved as relatively harmless sidebands separated from a dominant, collision-free carrier [28, 251, 252]. Although the effectiveness of these approaches is yet to be demonstrated at the sub-10^{-16} level, there is reason to expect that collisions could be prevented from undermining state-of-the-art uncertainty in lattice clocks with ^{88}Sr.

4.2.2 Hyper-Ramsey spectroscopy

The probe Stark shift problem is also shared by a handful of other species of optical clock, most notably the highly-forbidden electric-octupole $^2S_{1/2} \to {}^2F_{7/2}$ ^{171}Yb$^+$ clock [117, 137]. A conceptually straightforward approach to characterising the shift is exemplified by recent realisations of the ^{171}Yb$^+$ clock, where the unperturbed atomic resonance was extrapolated from interleaved Rabi-spectroscopy sequences of high and low probe intensity [91, 117]. Although careful extrapolation can be quite effective, achieving frequency uncertainty nearly 4 orders of magnitude smaller than the shift itself [91], the stability of the clock is deteriorated by the extrapolation process and the ultimate accuracy is limited by the precision with which the probe intensity ratio can be calibrated.

In order to reduce the burden of probe intensity control, tailored spectroscopy pulses were proposed which provide a central feature at $\nu \approx \nu_0$ whose frequency is unchanged by the light shift [255, 297, 298, 300]. The great potential of these "hyper-Ramsey" protocols was first illustrated on ^{171}Yb$^+$, where a 10^{-13} probe Stark shift was suppressed by more than four orders of magnitude below the 10^{-17} level [115]. As well as improving the clock's accuracy, hyper-Ramsey spectroscopy also helped to maintain a good clock stability: there is no need to undergo any kind of noise-amplifying extrapolation process, and the hyper-Ramsey frequency discriminant at $\nu \approx \nu_0$ is almost as steep as in standard Ramsey fringes (i.e. pretty steep!).

In this section we describe the theory and the practical implementation of hyper-Ramsey spectroscopy in our ^{88}Sr lattice clock. As well as being the first group to apply this type of spectroscopy in a lattice clock, we also invented and demonstrated a new "modified hyper-Ramsey" scheme which suppresses the probe Stark shift much more efficiently, even in the face of large variations in probe intensity. We wrote up

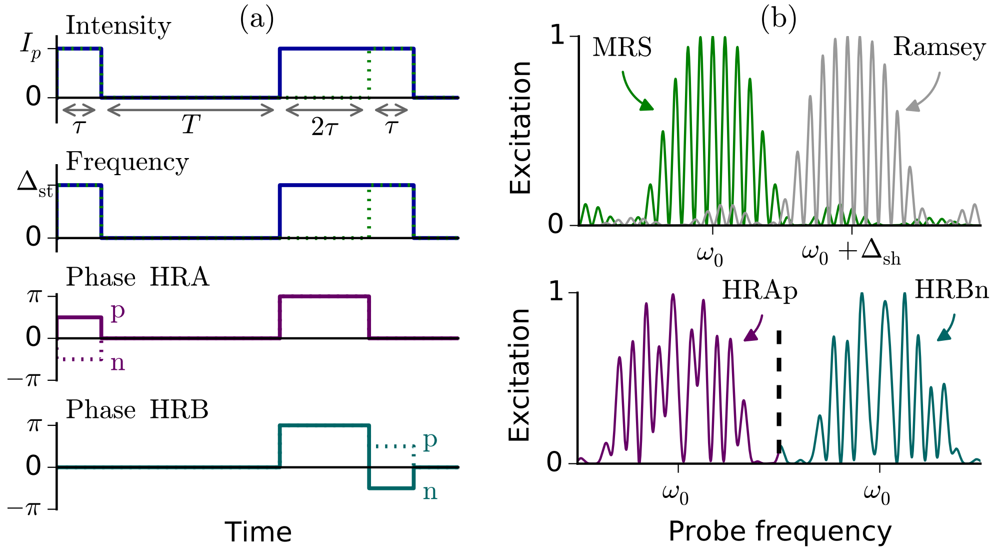

Figure 4.5: Pulse patterns (a) and excitation spectra (b) of the various spectroscopy methods considered in this paper. In the top half of (a), the blue traces indicate the common intensity and frequency patterns of hyper-Ramsey types A and B, while the dotted green traces show modified Ramsey spectroscopy (MRS). The spectra in (b) are calculated with $T = 4\tau$, $\Omega_0\tau = \pi/2$, $\Delta_{sh}/2\pi = 1.56/\tau$, and $\Delta_{st} = \Delta_{sh}$.

our results, released the paper onto ArXiv, and it was soon accepted for publication as a rapid communication in Physical Review A [110]. However, straight after our paper went onto ArXiv, we were informed that a similar scheme, dubbed "generalized hyper-Ramsey", had just been invented by Thomas Zanon-Willette at SYRTE – as it turns out, he was writing up a theoretical paper on the topic at roughly the same time we were [299]. In section 4.2.2.3 we present a unifying formalism for both 'modified' and 'generalized' hyper-Ramsey protocols, describing how the different spectroscopy patterns belong to a continuum of equally probe-shift-immune schemes.

4.2.2.1 Theory of modified hyper-Ramsey spectroscopy

Taking the Ramsey method of separated oscillatory fields [218] as the starting point for our discussion of probe-immune spectroscopy techniques, we observe two problems which emerge when a large probe shift is present: the envelope of the Ramsey fringes moves to follow the shift, and the distribution of fringes themselves becomes asymmetric around the peak (see figure 4.5). A solution is pointed out in [255], which we call "modified Ramsey" spectroscopy (MRS): if we apply a frequency step Δ_{st} to the probe laser to compensate exactly for the shift Δ_{sh}, then both the envelope and the Ramsey fringes are restored exactly back to the position of the perfect, unperturbed Ramsey spectrum. However, applying the correct frequency step requires us to know exactly what the light shift is, which is experimentally unrealistic. When the compensation step is slightly incorrect – i.e. when a residual uncompensated shift $\Delta = \Delta_{\text{sh}} - \Delta_{\text{st}}$ remains – the frequency of the central MRS fringe moves, feeding directly into a frequency shift of the locked clock.

To reduce sensitivity to Δ, it was proposed in [297] to insert an extra pulse into the Ramsey dark time. In detail, this proposal corresponds to the "hyper-Ramsey type A" spectroscopy (HRA) depicted in figure 4.5: using a constant laser intensity during the pulses, we expose the atoms to a pulse of length τ and phase $\pm\pi/2$, then wait for dark time T, and finally apply a pair of pulses of length 2τ, τ and phases π, 0 respectively. The intensity should ideally be chosen to give a first pulse area close to $\pi/2$ to maximize the frequency discriminant, but errors in pulse area have no significant effect on the probe shift suppression [297, 300]. The atomic clock is stabilized to the hyper-Ramsey feature as follows: Interleaving one sequence of HRA_{p} (phase $+\pi/2$) followed by one sequence of HRA_{n} (phase $-\pi/2$) we use the difference in excitation fractions as an error signal to feedback to the clock frequency. In the long term, the clock frequency will settle at the point where the two excitation probabilities P are the same, i.e. where $P_{\text{HRA}_{\text{p}}} = P_{\text{HRA}_{\text{n}}}$. Using the HRA spectroscopy pattern, not only is the locked frequency equal to the unperturbed ω_0 when we apply exactly the correct compensation step, but its linear dependence on small variations in Δ is also eliminated [297]. However, the HRA clock remains vulnerable to a residual cubic dependence on Δ (see figure 4.7), thus still requiring careful control of the compensation step Δ_{st} to within a small region around Δ_{sh} to avoid significant shifts.

An essential part of our modified scheme is the new hyper-Ramsey 'type B' (HRB) spectroscopy pattern depicted in figure 4.5, which is identical to HRA except that the $\pm\pi/2$ phase step is implemented in the last pulse instead of the first. As we show below, using a mix of both HRA and HRB spectroscopy we can ultimately eliminate all dependence of the locked clock on Δ, providing a clock lockpoint at exactly ω_0 regardless of errors in the compensation step.

In the rotating frame of the laser field $\mathbf{E}(t) = \mathbf{E}_0 \cos((\omega_L + \Delta_{\text{st}})t + \phi)$, we model the evolution of the atomic state during the probe pulses in the basis $|g\rangle = \binom{0}{1}$ and $|e\rangle = \binom{1}{0}$ using the propagator:

$$\hat{W}(t_p, \Omega_0, \Delta_p, \phi) = \begin{pmatrix} \cos(\frac{\Omega t_p}{2}) + i\frac{\Delta_p}{\Omega}\sin(\frac{\Omega t_p}{2}) & -ie^{-i\phi}\frac{\Omega_0}{\Omega}\sin(\frac{\Omega t_p}{2}) \\ -ie^{i\phi}\frac{\Omega_0}{\Omega}\sin(\frac{\Omega t_p}{2}) & \cos(\frac{\Omega t_p}{2}) - i\frac{\Delta_p}{\Omega}\sin(\frac{\Omega t_p}{2}) \end{pmatrix} \tag{4.25}$$

where we define a probe time t_p, a Rabi frequency $\Omega_0 = \mathbf{d} \cdot \mathbf{E}/\hbar$, a generalized Rabi frequency $\Omega = \sqrt{\Omega_0^2 + \Delta_p^2}$, and an effective laser detuning $\Delta_p = \omega_L - \omega_0 - \Delta$. The effects of spontaneous decay and decoherence are neglected.

Meanwhile, the propagator for the hyper-Ramsey dark time T is given by:

$$\hat{V}(T, \delta) = \begin{pmatrix} e^{i\frac{\delta T}{2}} & 0 \\ 0 & e^{-i\frac{\delta T}{2}} \end{pmatrix} \tag{4.26}$$

where $\delta = \omega_L - \omega_0$.

For the example of hyper-Ramsey spectroscopy of type A with a positive $\pi/2$ phase step, we can then write the total propagator for the sequence as:

$$\hat{U}_{\mathrm{HRA_p}} = \hat{W}(\tau, \Omega_0, \Delta_p, 0)\hat{W}(2\tau, \Omega_0, \Delta_p, \pi)\hat{V}(T, \delta)\hat{W}(\tau, \Omega_0, \Delta_p, \pi/2) \tag{4.27}$$

while propagators for the other three types of hyper-Ramsey look identical except for different phases in the first and last pulse.

We observe a set of important symmetries for the evolution during the pulses (equation 4.25):

$$\hat{W}(t_p, \Omega_0, \Delta_p, \phi)^\dagger = \hat{W}(t_p, \Omega_0, -\Delta_p, \phi + \pi) \tag{4.28}$$

$$\hat{W}\left(t_p, \Omega_0, \Delta_p, (0, \pi)\right)^T = \hat{W}\left(t_p, \Omega_0, \Delta_p, (0, \pi)\right) \tag{4.29}$$

$$\hat{W}\left(t_p, \Omega_0, \Delta_p, \pm\pi/2\right)^T = \hat{W}\left(t_p, \Omega_0, \Delta_p, \mp\pi/2\right) \tag{4.30}$$

If we consider hyper-Ramsey sequences where the laser is on resonance such that $\delta = 0$, $\Delta_p = \Delta$ and $\hat{V}(T, \delta) = \mathbb{I}$, we can use the symmetry eqns 4.28 - 4.30 to derive general relationships between the propagators for different types of hyper-Ramsey interrogation. For instance:

$$\begin{aligned} \hat{U}^T_{\delta=0,\mathrm{HRA_p}} &= \left(\hat{W}(\tau, \Omega_0, \Delta, 0)\hat{W}(2\tau, \Omega_0, \Delta, \pi)\hat{W}(\tau, \Omega_0, \Delta, \pi/2)\right)^T \\ &= \hat{W}(\tau, \Omega_0, \Delta, \pi/2)^T\hat{W}(2\tau, \Omega_0, \Delta, \pi)^T\hat{W}(\tau, \Omega_0, \Delta, 0)^T \\ &= \hat{W}\left(\tau, \Omega_0, \Delta, -\pi/2\right)\hat{W}(2\tau, \Omega_0, \Delta, \pi)\hat{W}(\tau, \Omega_0, \Delta, 0) \\ &= \hat{U}_{\delta=0,\mathrm{HRB_n}} \end{aligned} \tag{4.31}$$

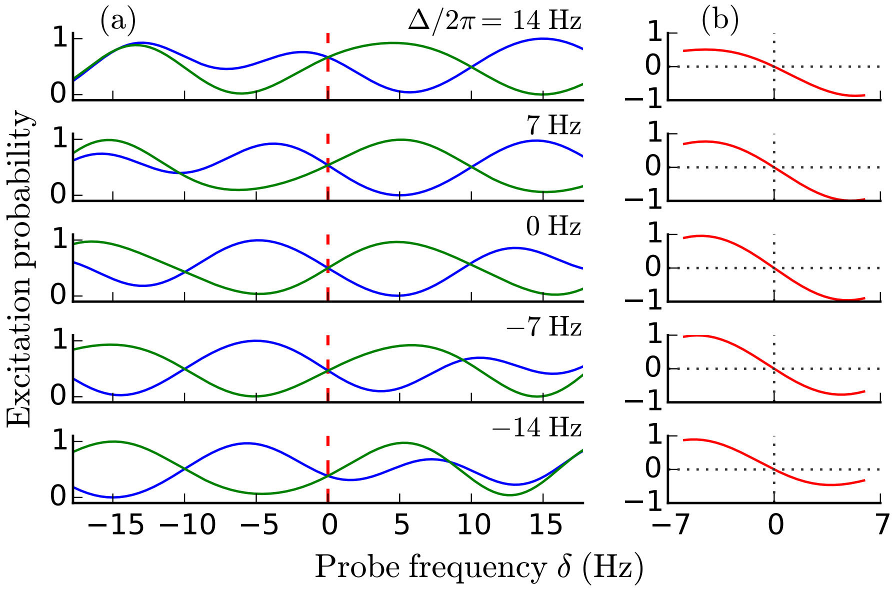

Figure 4.6: (a) Excitation spectra and (b) error signals $P_{\mathrm{HRA_p}} - P_{\mathrm{HRB_n}}$ around the central hyper-Ramsey fringe for different uncompensated shifts. In (a), the blue and green curves correspond to $\mathrm{HRA_p}$ and $\mathrm{HRB_n}$ sequences respectively and the red dashed lines highlight that the excitation probabilities always cross at $\delta = 0$. We use $\Omega_0\tau = \pi/2$, $\tau = 10\,\mathrm{ms}$, and $T = 50\,\mathrm{ms}$.

and similarly $\hat{U}^T_{\delta=0,\mathrm{HRA_n}} = \hat{U}_{\delta=0,\mathrm{HRB_p}}$.

Using the fact that all the total propagators must be unitary, these results imply equalities between the on-resonance excitation probabilities:

$$\begin{aligned} P_{\mathrm{HRA_p}} &= |\langle e|\, \hat{U}^T_{\delta=0,\mathrm{HRA_p}}\, |g\rangle\,|^2 \\ &= |\langle e|\, \hat{U}_{\delta=0,\mathrm{HRB_n}}\, |g\rangle\,|^2 = P_{\mathrm{HRB_n}} \end{aligned} \tag{4.32}$$

and similarly $P_{\mathrm{HRA_n}} = P_{\mathrm{HRB_p}}$.

The effect of the equality in equation 4.32 is illustrated in figures 4.6 and 4.7. Over a range of different uncompensated shifts Δ, we see that the hyper-Ramsey fringes always cross at the unperturbed atomic resonance $\omega_L = \omega_0$. This represents the critical result of our scheme: Using an error signal proportional to either $P_{\mathrm{HRA_p}} - P_{\mathrm{HRB_n}}$ or $P_{\mathrm{HRB_p}} - P_{\mathrm{HRA_n}}$ we realize a locked clock laser frequency at ω_0 which is immune to variations in Δ. Note also that this result is independent of pulse area $\Omega_0\tau$, meaning that exact realization of a $\pi/2$ pulse area is not needed.

However, although the lineshape crossing is robust against Δ, figure 4.6 also reveals that the slope at the crossing is reduced if the uncompensated shift is too large. If this reduced slope were allowed to persist, the stability of the locked clock would be compromised and its accuracy could be contaminated by asymmetries in the frequency discriminant around the zero-crossing. Therefore, even in the modified hyper-Ramsey scheme we still need to make sure that the shift compensation is roughly correct. It would be possible to interrupt the clock operation with a separate Rabi servo to evaluate the probe shift over time (as implemented in [115]), but another useful symmetry between the different types of hyper-Ramsey spectroscopy provides us with a method of servoing Δ towards 0 without degrading clock stability. In particular, we have a $\delta = 0$ evolution:

$$\begin{aligned}
\hat{U}^{\dagger}_{\delta=0,\mathrm{HRA_p}} &= \Big(\hat{W}(\tau,\Omega_0,\Delta,0)\hat{W}(2\tau,\Omega_0,\Delta,\pi)\hat{W}(\tau,\Omega_0,\Delta,\pi/2)\Big)^{\dagger} \\
&= \hat{W}(\tau,\Omega_0,\Delta,\pi/2)^{\dagger}\hat{W}(2\tau,\Omega_0,\Delta,\pi)^{\dagger}\hat{W}(\tau,\Omega_0,\Delta,0)^{\dagger} \\
&= \hat{W}\left(\tau,\Omega_0,-\Delta,-\pi/2\right)\hat{W}(2\tau,\Omega_0,-\Delta,0)\hat{W}(\tau,\Omega_0,-\Delta,\pi) \\
&= \hat{W}\left(\tau,\Omega_0,-\Delta,\pi/2\right)\hat{W}(2\tau,\Omega_0,-\Delta,\pi)\hat{W}(\tau,\Omega_0,-\Delta,0) \\
&= \hat{U}_{\delta=0,\mathrm{HRB_p}}(-\Delta) \qquad (4.33)
\end{aligned}$$

where in the penultimate step we have added a global phase π to all the pulses which will not affect the atom dynamics. This implies further equalities between on-resonance excitation probabilities:

$$P_{\mathrm{HRA_p}}(\Delta) = P_{\mathrm{HRB_n}}(\Delta) = P_{\mathrm{HRA_n}}(-\Delta) = P_{\mathrm{HRB_p}}(-\Delta) \qquad (4.34)$$

The results of this equality are illustrated in the inset of figure 4.8: the equilibrium excitation fraction changes with Δ in both the $\mathrm{HRA_p}$ & $\mathrm{HRB_n}$ and the $\mathrm{HRA_n}$ & $\mathrm{HRB_p}$ modified hyper-Ramsey locks, but these changes are equal and opposite. In the main part of the figure we take the difference between the equilibrium excitations fractions $1/2(P_{\mathrm{HRA_n}}+P_{\mathrm{HRB_p}})-1/2(P_{\mathrm{HRA_p}}+P_{\mathrm{HRB_n}})$ which can clearly be used as an error signal to steer the compensation step Δ_{st} towards Δ_{sh}.

4.2.2.2 Hyper-Ramsey spectroscopy on ^{88}Sr

Further to the theoretical models, we implement the modified hyper-Ramsey scheme on our ^{88}Sr lattice clock to gather data plotted in figures 4.7 and 4.8. For this data we employ a shortened version of the 'streak loading' technique described in section 2.3.2: the single-frequency red MOT lasts a total of 120 ms, and in the latter half of this time we ramp the 689 nm detuning by −200 kHz to allow the MOT to fall under

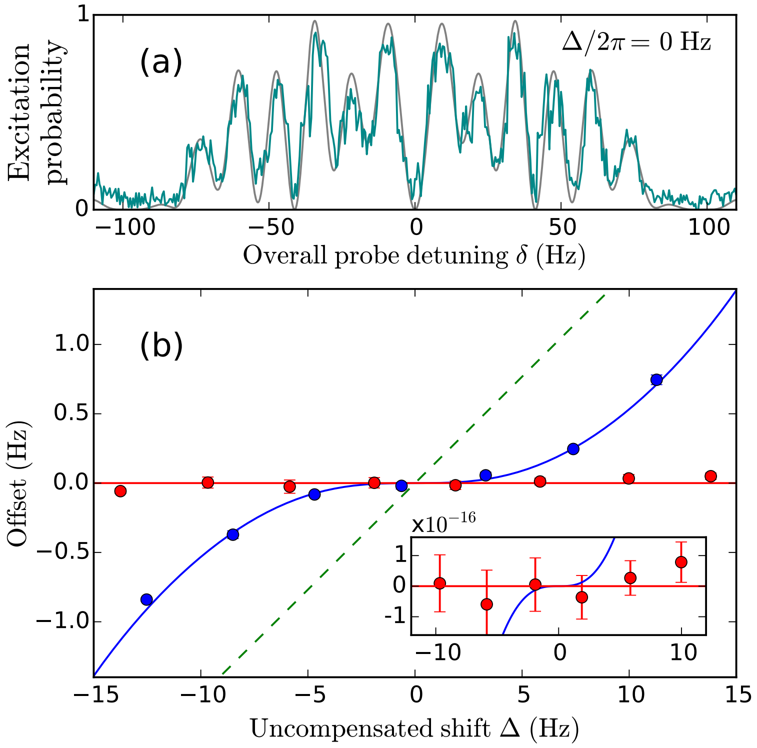

Figure 4.7: (a) Scan over the hyper-Ramsey feature in ^{88}Sr for $\tau = 10\,\text{ms}$, $T = 50\,\text{ms}$, and $\Omega_0\tau \approx \pi/2$, and with $\phi = 0$ in the first and last pulse. The theoretical model is overlaid in gray with no fitting parameters used. (b) Modelled and measured residual Stark shifts for the different spectroscopy methods: The HRA$_\text{p}$ & HRA$_\text{n}$ 'standard' hyper-Ramsey lock (blue) shows good suppression compared with modified Ramsey (dashed green), but the HRA$_\text{n}$ & HRB$_\text{p}$ modified hyper-Ramsey (red) is better. *Inset:* Enlarged view showing the residual lock offset in fractional frequency units.

gravity in a controlled manner. At the end of this loading stage, the atoms are spread along a 300 µm streak covering around 700 lattice sites. Before interrogating the clock transition, we ramp the lattice depth from 11 µK to around 2 µK in 30 ms, hold for 30 ms, and then ramp back up to 11 µK in another 30 ms. This final atom-filtering stage ensures that the hotter atoms are discarded and gives a stable method to load at low atomic density.

We interrogate the magnetically-induced 698 nm ^{1}S$_0$ − ^{3}P$_0$ clock transition in ^{88}Sr using a mixing field $|B| \approx 2.5\,\text{mT}$. We focus 2.7 mW of interrogating light to a waist of approximately 250 µm at the position of the trapped atoms. In order to compensate for AOM chirp and fibre phase noise the clock delivery is actively phase-stabilized, with constant compensation of the final 'switching' AOM achieved by retro-reflecting the zeroth-order beam (see section 4.1.6).

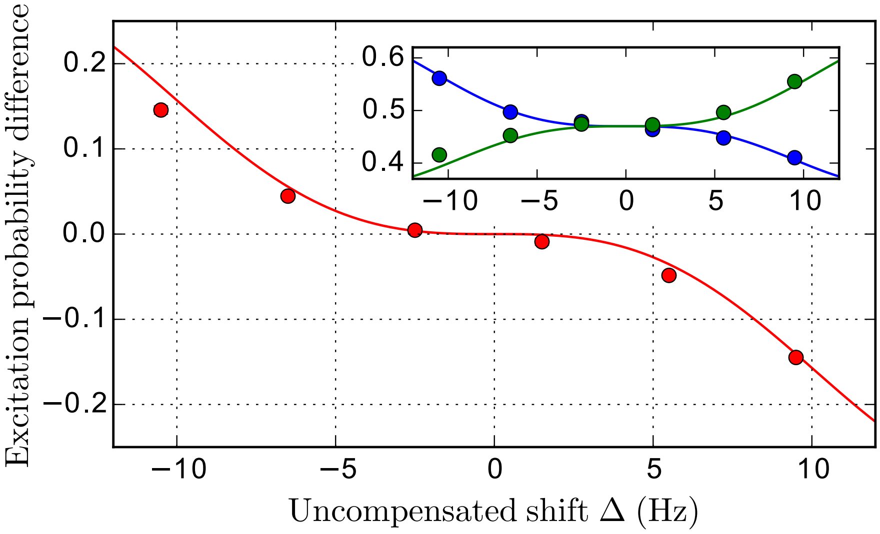

Figure 4.8: *Inset:* The dependencies of equilibrium $\delta = 0$ excitation fraction on uncompensated shift for both modified hyper-Ramsey locks: $\mathrm{HRA_p}$ & $\mathrm{HRB_n}$ (green) and $\mathrm{HRA_n}$ & $\mathrm{HRB_p}$ (blue). *Main:* Excitation difference to be used for steering the compensation step towards $\Delta = 0$. The model is scaled by 0.94 to match the experimental fringe contrast.

For the data in figure 4.7 we operate with 4 active clock servos. Servo 1 is the reference, using a modified hyper-Ramsey $\mathrm{HRA_p}$ & $\mathrm{HRB_n}$ sequence with a compensation step of $\Delta_{\mathrm{st}} = 80\,\mathrm{Hz}$ close to the slowly-varying probe Stark shift $79.6 < \Delta_{\mathrm{sh}}/2\pi < 80.7\,\mathrm{Hz}$. Servo 2 then uses one of the two types of hyper-Ramsey sequence investigated in the figure, applying various different compensation steps $66 < \Delta_{\mathrm{st}}/2\pi < 94\,\mathrm{Hz}$ to map out the residual Stark shift. The y-position of plotted points is given by the frequency offset between servos 1 and 2. Meanwhile, servos 3 and 4 utilise Rabi interrogation as a live monitor the Stark shift. They are set to run in only 4 out of 16 cycles to avoid degrading the stability of the hyper-Ramsey servos. Servo 3 uses the same intensity and magnetic field as the hyper-Ramsey servos and is therefore shifted by Δ_{sh} from servo 1, allowing us to calculate the x-position of the plotted points in the figure. The statistical uncertainty in the frequency offset between servo 3 and servo 1 is always less than 0.1 Hz, so the horizontal error bars are smaller than the data points. Servo 4 uses a lower intensity but the same field as servo 3, allowing us to extrapolate an independent (but less accurate) measure of Δ_{sh}. For the data in figure 4.8 we use the same set of 4 servos except that the compensation step in servo 1 is no longer set 80 Hz but is instead scanned in tandem with that of servo 2.

To conclude, the data in figure 4.7 verifies our most important result: within the 1×10^{-16} measurement statistics the modified hyper-Ramsey scheme proves com-

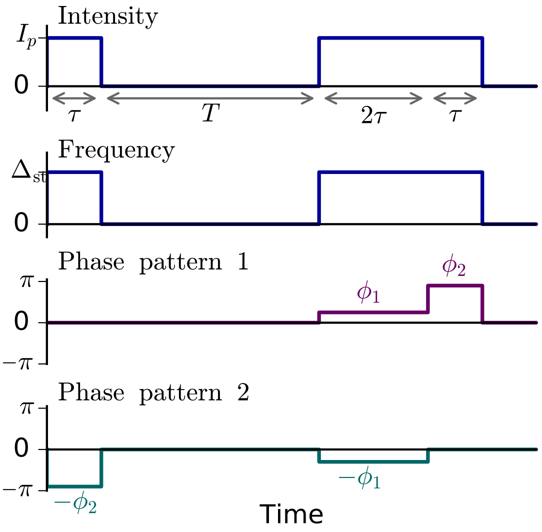

	ϕ_1	ϕ_2
Modified HR	π	$-\frac{\pi}{2}$
Modified HR	π	$\frac{\pi}{2}$
'Generalized' HR	$\frac{\pi}{4}$	0
'Generalized' HR	$\frac{3\pi}{4}$	0
Generalised HR	ϕ_1	$2\phi_1 \pm \frac{\pi}{2}$

Figure 4.9 & Table 4.2: A schematic showing generalised pairs of hyper-Ramsey patterns GHR_1 and GHR_2 whose lineshapes always cross at $\delta = 0$ (see equation 4.36), thus creating a lockpoint which is immune to probe Stark shifts. *Table:* Specific phases ϕ_1, ϕ_2 for previously published hyper-Ramsey schemes [110, 299], plus the most general phase combination which still yields maximum discriminant at $\delta = 0$.

pletely immune to the 2×10^{-13} probe Stark shift over a large range of compensation steps. Furthermore, the data in figure 4.8 verifies the expected antisymmetric dependence of equilibrium excitation fraction on Δ, creating an easily-accessible error signal for constantly steering the compensation step towards $\Delta_{\mathrm{st}} \approx \Delta_{\mathrm{sh}}$.

4.2.2.3 Generalised hyper-Ramsey pulses

As it turns out, the modified hyper-Ramsey protocol presented above is just one among a continuum of pulse patterns, each of which can eliminate the probe Stark shift equally efficiently. To prove this, we observe that equation 4.25, governing the atomic evolution during the probe pulses, actually follows a more general symmetry than we originally pointed out in equations 4.28 − 4.30. In particular, we notice that

$$\hat{W}(t_p, \Omega_0, \Delta_p, \phi)^T = \hat{W}(t_p, \Omega_0, \Delta_p, -\phi) \tag{4.35}$$

implying that the propagators for the two generalised hyper-Ramsey phase patterns GHR_1 and GHR_2 depicted in figure 4.9 follow the identity

$$
\begin{aligned}
\hat{U}^{T}_{\delta=0,\mathrm{GHR}_1} &= \Big(\hat{W}(\tau,\Omega_0,\Delta,\phi_2)\hat{W}(2\tau,\Omega_0,\Delta,\phi_1)\hat{W}(\tau,\Omega_0,\Delta,0)\Big)^{T} \\
&= \hat{W}\left(\tau,\Omega_0,\Delta,0\right)\hat{W}(2\tau,\Omega_0,\Delta,-\phi_1)\hat{W}(\tau,\Omega_0,\Delta,-\phi_2) \\
&= \hat{U}_{\delta=0,\mathrm{GHR}_2}
\end{aligned} \tag{4.36}
$$

Using an identical argument as in section 4.2.2.1, equation 4.36 proves that the lineshapes of the pair of generalised hyper-Ramsey pulses always cross at $\delta = 0$ regardless of our choice of ϕ_1 and ϕ_2, thus realising a frequency lockpoint which is immune from variations in the probe Stark shift.

Another important consideration for clock operation is the frequency discriminant at $\delta = 0$. For a given ϕ_1, we find that the optimum final phase step is given by $\phi_2 = 2\phi_1 + (2k+1)\frac{\pi}{2}$ for integer k, though since integer multiples of 2π have no effect, we can pick the two cases $\phi_2 = 2\phi_1 \pm \frac{\pi}{2}$ without loss of generality. The appropriateness of these choices of ϕ_2 can be confirmed by following the evolution of the atomic state on the Bloch sphere – in particular, we find that the clock lockpoint for $\phi_2 = 2\phi_1 \pm \frac{\pi}{2}$ lies exactly at the steepest point on the side of the fringe, i.e. where $P_{\mathrm{GHR}}(\delta = 0) = 0.5$, assuming otherwise ideal pulse parameters $\Omega_0\tau = \pi/2$ and $\Delta = 0$.

The degree of freedom we have in choosing ϕ_1 and ϕ_2 may seem unnecessary – the specific phase choices in the modified and ‘generalized’ hyper-Ramsey protocols already completely eliminated the probe Stark shift – but it does have a significant advantage: since the different phase combinations have different dependences on parasitic effects such as motional heating and laser decoherence (see the following section), we are free to choose the most accurate phase steps (or indeed combinations of phase steps) for which these parasitic offsets are minimised.

4.2.2.4 Residual probe shifts

In theory, the modified hyper-Ramsey scheme (and its generalisation) should exhibit no probe Stark shift at all. However, in practice these schemes are susceptible to a handful of potentially dangerous residual probe-induced offsets:

- Asymmetry in the error signal can be sampled if the clock laser occasionally deviates a long way from $\delta = 0$, ultimately causing slower net frequency corrections in one direction that in the other. This can be minimised by choosing a conservative pulse time much shorter than the clock laser coherence time, and by compensating the Stark shift near to $\Delta_{\mathrm{st}} \approx \Delta_{\mathrm{sh}}$ to minimise lineshape asymmetry (see figure 4.6).

- Motional heating of the atom(s) causes the carrier Rabi frequency to decrease from the beginning to the end of the pulse sequence (higher motional n leads to lower Ω_0 as seen in section 2.3.3.2), causing a net deviation from the ideal hyper-Ramsey pulse areas. This is particularly a problem for ion clocks, which tend to have substantial heating rates of order 100 quanta per second. Carrying out numerical models of an adjusted hyper-Ramsey sequence including a linear ramp in Ω_0, we observe a sizeable effect for RF-trapped ions, but not so much for lattice-trapped atoms.

- As pointed out in [296], decoherence of the clock laser (or, equivalently, of the atomic transition) modifies the atomic evolution during the hyper-Ramsey pulse sequences, introducing a subtle systematic offset to the modified hyper-Ramsey lock point.

- Further deviations from the intended hyper-Ramsey pulse sequence can be caused by AOM-induced phase transients and phase chirps in the optical probe beam (see section 4.1.6).

In the optical lattice clock we expect these effects to be controllable well below the 10^{-18} level, though they will all need careful characterisation before we can assign a total systematic uncertainty for the residual probe Stark shift.

In reality, the biggest residual probe-related shift in our ^{88}Sr lattice clock is the quadratic Zeeman shift, which is around 150 Hz for our operating 2.5 mT bias field. Unfortunately, it is not possible for us to eliminate the quadratic Zeeman shift using the hyper-Ramsy techniques, simply because it is impossible to extinguish the magnetic field quickly enough during the hyper-Ramsey dark time – although we can switch the coil current in around 200 µs, the steel chamber (and copper gaskets) retain eddy currents which decay much more slowly with a time constant on the order of 1.5 ms. Since these eddy currents are difficult to measure with the requisite precision, and there's no guarantee even of their shot-to-shot repeatability, we find it prudent to avoid changing the value of our magnetic field at any time during the spectroscopy pulse.

Since we can't eliminate the quadratic Zeeman shift using hyper-Ramsey techniques, the best remaining option is simply to measure it. To do this, we intend to map out the clock frequency over a large range of bias magnetic fields so that the quadratic shift coefficient can be determined, using the ^{87}Sr stretched-state splitting as a reference for field calibration (see section 4.1.2). Judging by our measured

~ 20 nT magnetic field stability, and our ability to apply mixing fields up to 2.5 mT, a coefficient measurement at roughly the 1×10^{-4} level should be accessible in our current system. Combined with some tighter focusing of the probe beam to reduce the waist from 250 µm to around 100 µm, so that the mixing field can be reduced to more like 1 mT while still driving the same Rabi frequency, we ought to be able to reach a total probe-related uncertainty below the 10^{-17} level in our ^{88}Sr lattice clock. With future upgrades to the local oscillator stability, this shift uncertainty could be further reduced by extending beyond the comparatively short 90 ms probe times used throughout section 4.2.2.2.

4.2.3 Multi-photon lattice clocks

As we have seen, magnetically-induced spectroscopy is vulnerable to significant systematic frequency shifts from the probe Stark and quadratic Zeeman effects. These deficiencies have spurred the pursuit of alternative methods to excite the clock transition in bosonic lattice clocks, leading to various proposals to use two or more optical photons to drive the $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_0$ line indirectly [113, 208, 233, 254, 298, 301]. Importantly, these multi-photon techniques circumvent the need for a bias magnetic field, thus all but eliminating the Zeeman shift, even allowing for background field fluctuations of several µT. However, none of the multi-photon excitation routes are free from probe Stark shifts: for example, to create a resolvable electromagnetically-induced transparency feature on the $^1\mathrm{S}_0 \xrightarrow{E1} {}^1\mathrm{P}_1 \xrightarrow{M1} {}^3\mathrm{P}_0$ transition requires a strong, shift-inducing dressing laser at 1354 nm to drive the weak magnetic dipole transition [233, 298, 301]. Nonetheless, as we just saw in section 4.2.2, the probe Stark shift (unlike the Zeeman shift) can be efficiently suppressed using the technique of hyper-Ramsey spectroscopy.

In this section, we describe two new, relatively low-shift excitation routes for bosonic lattice clocks, based respectively on two ($E1M1$) and three ($E1E1E1$) photons. Although the two photon technique requires fewer probe beams, it also induces a larger Stark shift and needs a rather awkward probe wavelength of 4941 nm to drive the Sr clock transition. (On the other hand, two-photon clocks with Hg and Yb look quite attractive). The three-photon technique is conceptually a little more complicated, but we will show that it actually isn't too daunting to implement experimentally. The apparatus for three-photon spectroscopy is currently under construction in our lab, and we hope that it could put sub-10^{-18} probe-related inaccuracies within reach for our ^{88}Sr lattice clock.

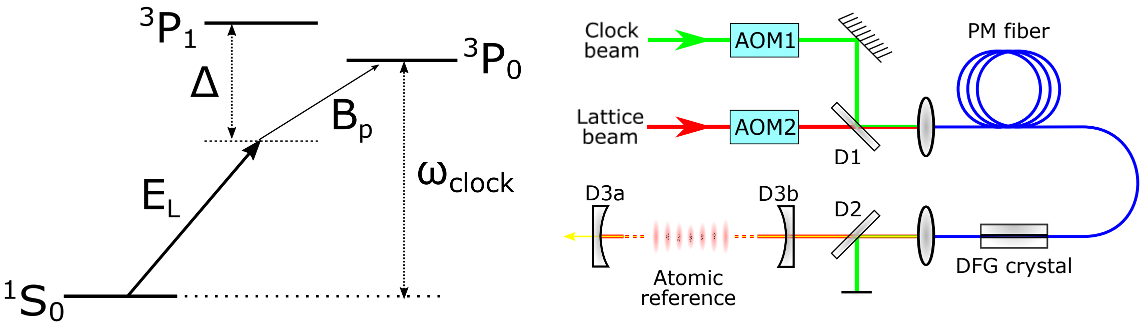

Figure 4.10: *Left:* Two-photon Raman excitation scheme in the three-level approximation. *Right:* A possible experimental realisation of the two-photon scheme. Dichroic D1 is used to mix the clock and lattice beams, which are sent through a waveguide non-linear crystal to generate the difference frequency (i.e. the probe beam). Dichroic D2 splits off the clock light from the lattice and probe beams to avoid unnecessary Stark shifts, and D3a and D3b form a lattice enhancement cavity. The AOMs perform duties of phase and intensity stabilization.

4.2.3.1 Two photons

We now describe our new two-photon $E1M1$ excitation route for bosonic optical lattice clocks. To minimize Stark shifts we propose to use the magic-wavelength lattice trap as one part of the probe. As depicted in figure 4.10, the other photon must be supplied from a supplementary beam at the frequency difference between the clock transition and the lattice laser[1].

The two-photon excitation rate can be calculated straightforwardly using first-order time-independent perturbation theory. Even though the polarisation dependence turns out to be pretty trivial, we include in our treatment all the $m = -1, 0, 1$ states within the $^3\mathrm{P}_1$ manifold – this full treatment lets us maintain a common language with our later section on the three-photon clocks (section 4.2.3.2). For now, we ignore auxiliary states other than $^3\mathrm{P}_1$.

The rest-frame atomic Hamiltonian in the presence of the lattice and probe fields can be written in the dipole approximation as:

$$\hat{H} = \sum_{\psi} E_\psi \left|\psi\right\rangle \left\langle\psi\right| - \hat{\mathbf{d}} \cdot \mathbf{E}_p(t) - \hat{\boldsymbol{\mu}} \cdot \mathbf{B}_L(t) \tag{4.37}$$

with fields $\mathbf{E}_L(t) = \mathbf{E}_L^{(+)} e^{-i\omega_L t} + \mathbf{E}_L^{(-)} e^{i\omega_L t}$ and $\mathbf{B}_p(t) = \mathbf{B}_p^{(+)} e^{-i\omega_p t} + \mathbf{B}_p^{(-)} e^{i\omega_p t}$

---

[1] A similar scheme was briefly mentioned in another paper [7], which explored the viability of two-photon clocks based on hot alkaline-earth vapours. There, the focus was on the use of two excitation photons of the same frequency; here, we remedy some minor errors in that treatment and we provide much more detailed calculations in the context of magically-excited optical lattice clocks.

which, when converted into the laser rotating frame, and discarding counter-rotating terms, gives a total Hamiltonian $\hat{H}^{(R)} = \hat{H}_0^{(R)} + \hat{H}_1^{(R)}$:

$$\hat{H}_0^{(R)} = -\left(\hbar\Delta \sum_m \left|{}^3P_1, m\right\rangle \left\langle {}^3P_1, m\right|\right) - \hbar\Delta_{\text{clock}} \left|{}^3P_0\right\rangle \left\langle {}^3P_0\right| \tag{4.38}$$

$$\begin{aligned}\hat{H}_1^{(R)} = \sum_m \Bigg[&- \left\langle {}^1S_0\right| \hat{\mathbf{d}} \cdot \mathbf{E}_L^{(-)} \left|{}^3P_1, m\right\rangle \left|{}^1S_0\right\rangle \left\langle {}^3P_1, m\right| \\ &- \left\langle {}^3P_1, m\right| \hat{\boldsymbol{\mu}} \cdot \mathbf{B}_p^{(+)} \left|{}^3P_0\right\rangle \left|{}^3P_1, m\right\rangle \left\langle {}^3P_0\right| \Bigg] + \text{h.c.}\end{aligned} \tag{4.39}$$

where Δ is the detuning of the lattice laser from the $\left|{}^3P_1\right\rangle$ manifold, and $\Delta_{\text{clock}} = \omega_0 - \omega_L - \omega_p$ is the detuning of the sum of the two probing photons from the clock transition.

The lattice electric field induces a first-order perturbation to the ground ${}^1\mathrm{S}_0$ state, yielding a dressed state:

$$\left\langle {}^1S_0'\right| = \left\langle {}^1S_0\right| - \frac{1}{\hbar\Delta} \sum_m \left\langle {}^1S_0\right| \hat{\mathbf{d}} \cdot \mathbf{E}_L^{(-)} \left|{}^3P_1, m\right\rangle \left\langle {}^3P_1, m\right| \tag{4.40}$$

allowing the probe field to drive a two-photon Rabi frequency given by[1]:

$$\begin{aligned}\Omega_{2\gamma} &= -\frac{2\left\langle {}^1S_0'\right| \hat{\boldsymbol{\mu}} \cdot \mathbf{B}_p^{(+)} \left|{}^3P_0\right\rangle}{\hbar} \\ &= \frac{2}{\hbar\Delta} \sum_m \left\langle {}^1S_0\right| \hat{\mathbf{d}} \cdot \mathbf{E}_L^{(-)} \left|{}^3P_1, m\right\rangle \left\langle {}^3P_1, m\right| \hat{\boldsymbol{\mu}} \cdot \mathbf{B}_p^{(+)} \left|{}^3P_0\right\rangle \\ &= \frac{2}{\hbar\Delta} \sum_{m,q_1,q_2} (-1)^{q_1+q_2} (\mathbf{E}_L^{(-)})_{-q_1} (\mathbf{B}_p^{(+)})_{-q_2} \left\langle {}^1S_0\right| \hat{d}_{q_1} \left|{}^3P_1, m\right\rangle \left\langle {}^3P_1, m\right| \hat{\mu}_{q_2} \left|{}^3P_0\right\rangle \\ &= \frac{2}{\hbar\Delta} \sum_m (-1)^{2m} (\mathbf{E}_L^{(-)})_m (\mathbf{B}_p^{(+)})_{-m} \frac{(-1)^{1+m}}{\sqrt{3}} \left\langle {}^1S_0\right| |\hat{\mathbf{d}}| \left|{}^3P_1\right\rangle \frac{1}{\sqrt{3}} \left\langle {}^3P_0\right| |\hat{\boldsymbol{\mu}}| \left|{}^3P_1\right\rangle \\ &= -\frac{2\left\langle {}^1S_0\right| |\hat{\mathbf{d}}| \left|{}^3P_1\right\rangle \left\langle {}^3P_0\right| |\hat{\boldsymbol{\mu}}| \left|{}^3P_1\right\rangle}{3\hbar\Delta} \sum_m (-1)^m (\mathbf{E}_L^{(-)})_m (\mathbf{B}_p^{(+)})_{-m}\end{aligned} \tag{4.41}$$

where q_1 and q_2 index the spherical-basis components of the $\mathbf{E}_L$ and $\mathbf{B}_p$ vector fields, and in the penultimate line we have applied the Wigner-Eckart theorem with appropriate Clebsch-Gordan coefficients. Once all the prefactors are collected together, the two-photon Rabi frequency can be rewritten in the simple form:

---

[1] We follow the conventions used in [242] for reduced dipole matrix elements. In particular, the relevant reduced magnetic moment is approximately given by $\left\langle {}^3P_0\right| |\hat{\boldsymbol{\mu}}| \left|{}^3P_1\right\rangle = \sqrt{2}\mu_B$, though from any given m state the magnitude of the matrix element is actually $\left|\left\langle {}^3P_0\right| \hat{\mu}_{-m} \left|{}^3P_1, m\right\rangle\right| = \sqrt{2/3}\mu_B$.

$$\Omega_{2\gamma} = \alpha\sqrt{I_L I_p}\,\hat{\mathbf{e}}_L \cdot \hat{\mathbf{b}}_p^* \tag{4.42}$$

where α is a measure of the induced Rabi frequency per unit of each field intensity, where $\hat{\mathbf{e}}_L$ is the unit lattice polarization vector, related to the oscillating electric field by $\mathbf{E}^{(-)} = \frac{|\mathbf{E}_L|}{2}\hat{\mathbf{e}}_L$, and where $\hat{\mathbf{b}}_p = \hat{\mathbf{k}}_p \times \hat{\mathbf{e}}_p$ is the unit probe magnetic field vector, related to the oscillating magnetic field by $\mathbf{B}^{(-)} = \frac{|\mathbf{B}_p|}{2}\hat{\mathbf{b}}_p$. Note that optimum conditions are satisfied when we choose a linear lattice polarisation parallel to the probe magnetic field vector; this can be realised using the straightforward geometry of a 1D lattice with a copropagating, orthogonally-polarized probe beam.

Moving beyond the three-level approximation, small adjustments to α must be taken into account. The largest correction is from the alternative excitation path ${}^1\mathrm{S}_0 \xrightarrow{E1} {}^1\mathrm{P}_1 \xrightarrow{M1} {}^3\mathrm{P}_0$, which contributes in the opposite phase to the dominant ${}^3\mathrm{P}_1$ route due to the opposing signs of the spin-orbit admixture between ${}^1\mathrm{P}_1$ and ${}^3\mathrm{P}_1$ [169]. Other $E1M1$ routes are only possible through configuration-mixing, and are further suppressed by their large detuning Δ, so they can safely be ignored.

A further contribution to $\Omega_{2\gamma}$ comes from the lattice magnetic field $\mathbf{B}_L$ plus the probe electric field $\mathbf{E}_p$, which also resonantly couples the clock states via ${}^3\mathrm{P}_1$ and ${}^1\mathrm{P}_1$. However, in the particular case of the 1D optical lattice, the atoms are trapped in the electric-field antinodes where $\mathbf{B}_L \approx 0$. Thus we will only get an effect from the lattice magnetic field if there is a significant running-wave component to the lattice laser beam. For completeness, we provide results for both standing-wave and running-wave configurations at the magic wavelength.

Of critical importance to optical clocks is the probe-induced frequency shift to the clock transition arising from off-resonant coupling to other atomic states. Since the lattice is operated at the magic wavelength, its net Stark shift contribution can in principle be maintained at zero. However, since a high lattice intensity is desired to drive the clock transition quickly, the hyperpolarisability becomes an important consideration. Thus, we can write the leading shifts as:

$$\Delta_{clock} = \kappa_p^{(1)} I_p + \kappa_L^{(2)}(\mathbf{e}_L) I_L^2 \tag{4.43}$$

where $\kappa_p^{(1)}$ and $\kappa_L^{(2)}(\mathbf{e}_L)$ are the shift coefficients due respectively to the differential polarisability at the probe wavelength and the differential hyperpolarisability at the magic wavelength. Since it is determined by two-photon transitions with varying angular-momentum selection rules, the hyperpolarisability in general depends on the

Table 4.3: A comparison of properties for different atomic species.

Atom	λ_L (nm)	λ_p (nm)	$\kappa_p^{(1)}$ $\left(\frac{\mathrm{Hz}}{\mathrm{W\,cm^{-2}}}\right)$	$\kappa_L^{(2)}$ $\left(\frac{\mathrm{Hz}}{(\mathrm{MW\,cm^{-2}})^2}\right)$	β $\left(\frac{\mathrm{s}^{-1}}{\mathrm{MW\,cm^{-2}}}\right)$	α_s	α_r $\left(\frac{\mathrm{Hz}}{\sqrt{\mathrm{W\,cm^{-2}}}\sqrt{\mathrm{MW\,cm^{-2}}}}\right)$
Sr	813	4941	-15.9	6.5 [145]	2.1	2.0	2.2
Yb	759	2428	-7.6	9.3 [16]	0.73	2.7	3.3
Hg	363	994	-1.3 [295]	0.10 [98]	0.26	0.86	1.1

lattice polarization $\hat{\mathbf{e}}_L$ (see section 4.1.3.2); here, we just give coefficients for the most common example of linear polarisation.

Another property desired in optical clocks is a long coherence time, so that a narrow line can be resolved to increase frequency sensitivity and ultimately to improve clock stability. This consideration places an additional upper limit on the feasible lattice intensity, since the rate of Rayleigh scattering increases as $\Gamma_L = \beta I_L$.

We list the important coefficients α, β and κ for Sr, Yb, and Hg in table 4.3, along with the relevant lattice and probe wavelengths. For the sake of clarity and due to the wider availability of data, we use values for the case of linear lattice polarization. However, an interesting possibility for Yb would be to use a magic ellipticity where the net hyperpolarisability shift is zero, allowing precision spectroscopy in very deep traps [256].

Considering only the hyperpolarisability and the scattering rate, it is tempting to conclude that the use of very high lattice intensities $\sim \mathrm{MW\,cm^{-2}}$ would be preferable in order to alleviate the probe Stark shift. However, as we explored in section 4.1.3, we need to be careful not to exacerbate the long series of residual non-linear Stark shifts proportional to $I_L^{1/2}$, I_L, $I_L^{3/2}$, and I_L^2. In very deep lattice traps, simultaneous characterisation of all these non-linearities to the 10^{-18} level would be quite difficult. In addition, collisional shifts and decoherence will also increase with trap depth due to tighter confinement of the atoms to a smaller volume [161], though these effects could be eliminated in an underfilled higher-dimensional lattices [5]. A tempting prospect would be to exploit hyper-Ramsey spectroscopy not just to eliminate the probe Stark shift, but also to suppress the residual lattice shifts, perhaps by ramping down the lattice intensity to a much lower value during the hyper-Ramsey dark time. However, manipulation of the trap depth during the probing sequence would risk changing the atomic motional state, thus critically contaminating the relative hyper-Ramsey pulse areas and potentially reintroducing parasitic Stark shifts from the probe beam (see section 4.2.2.4).

Nonetheless, even without hyper-Ramsey elimination of the lattice shifts, the two-photon scheme does look viable as the basis for a highly accurate clock with Yb or Hg – for typical standard peak lattice intensities of $140\,\mathrm{kW/cm^2}$ [211] and $300\,\mathrm{kW/cm^2}$ [268] in Yb and Hg respectively, a Rabi frequency of 3.8 Hz (4.1 Hz) can be driven with a probe shift of 100 Hz – perfectly manageable using hyper-Ramsey techniques [117]. For reference, these Rabi frequencies are only a factor of two smaller than in the $^{171}\mathrm{Yb}^+$ octupole clock, assuming the same probe Stark shift. Sadly, the two-photon Rabi frequency in Sr is around 5 times smaller than in Yb or Hg for typical lattice depths, so we conclude that the two-photon Sr clock doesn't sound that promising.

An important practical issue for multi-photon schemes is how to realize them in experiment. The challenge is to stabilise the sum frequency $\nu_L + \nu_p$ to the sub-Hz level required to perform precision clock spectroscopy. In earlier proposals using two photons, it was suggested that one could sum the photons $\omega_L + \omega_p$ in a non-linear crystal, and then stabilise the output with an optical phase-locked loop against a clock local oscillator running at ω_{21} [233]. However, if enough power is available at the clock wavelength ($\gtrsim 10\,\mathrm{mW}$) we suggest the alternative implementation depicted in fig. 4.10: A non-linear crystal is used to take the difference frequency of the lattice beam with the existing clock laser beam near the position of the atoms, thus generating phase-stable light at the desired probe frequency $\nu_p = \nu_{21} - \nu_L$. Exploiting the high efficiency of commercially-available waveguide PPLN crystals, and utilizing the large amount of lattice power available, it may be possible to generate sufficient probe light ($\sim$ mW) without the need for a new laser at the probe wavelength, at least for Yb or Sr. If necessary an injection-locked slave diode laser could be installed as a simple means of probe amplification.

For Hg, the deep UV wavelengths make efficient non-linear frequency conversion of the clock and lattice beams more difficult. However, there is viable scheme with non-linear crystals which could avoid the need to generate any beam at the awkward clock wavelength of 265.6 nm: instead, one could build an ultrastable 521 nm local oscillator at $\omega_{\mathrm{LO}} = \omega_{\mathrm{clock}}/2$, and send this laser into two non-linear crystals with the lattice and probe beams to generate two 1141 nm beams at $\omega_{\mathrm{diff1}} = \omega_L - \omega_{\mathrm{LO}}$ and $\omega_{\mathrm{diff2}} = \omega_{\mathrm{LO}} - \omega_p$. If we phase lock the optical beat between these two output beams so that $\omega_{\mathrm{diff1}} \approx \omega_{\mathrm{diff2}}$, actuating on the 994 nm probe laser, we can ultimately impose the necessary coherence on the frequency sum $\omega_L + \omega_p \approx 2\omega_{\mathrm{LO}}$. This all may sound like a heroic effort, but I'd rather work with two waveguide differencing crystals plus

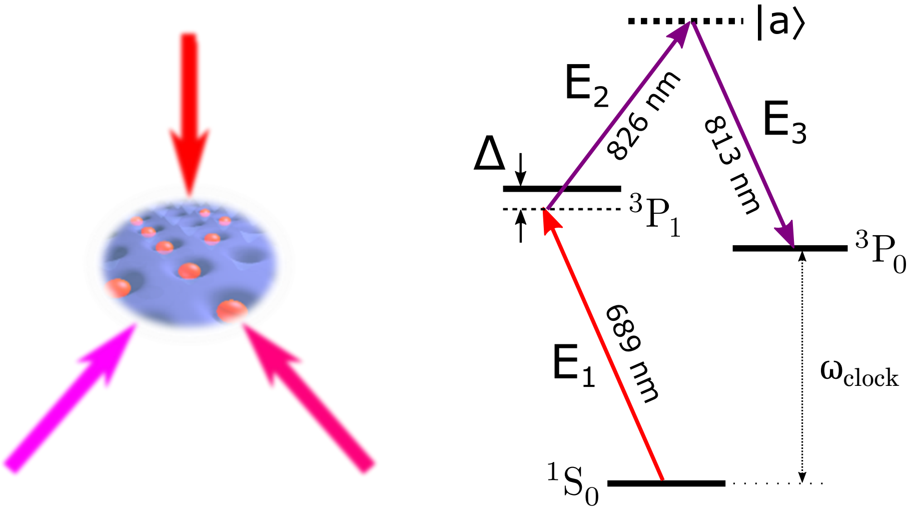

Figure 4.11: *Left:* Illustration of recoil-free three-photon spectroscopy in a 3D optical lattice. *Right:* The three-photon excitation route at a glance.

a phase lock than to have to build and maintain a resonant SHG cavity to generate deep-UV 265.6 nm light.

4.2.3.2 Three photons

We depict the three-photon excitation scheme in figure 4.11. Once again, we absorb one photon at 813 nm, potentially from the lattice trap, in order to minimise probe Stark shifts. A second probe beam at 826 nm also induces a relatively small Stark shift due to its proximity to the magic wavelength. The third photon is supplied by our 2^{nd}-stage cooling laser at 689 nm, but is detuned from the cooling transition by on the order of 10 MHz in order to avoid significant scattering of 689 nm photons.

Since the 813 nm probe beam is far off resonance from any transition, there is a long series of auxiliary states $|a\rangle$ which mediate the three-photon excitation. Notably, the dipole selection rules ensure that all the auxiliary states must possess total electronic angular momentum $J_\alpha = 1$. The $5s6s$ $^3\mathrm{S}_1$ manifold is the nearest-detuned strongly-allowed auxiliary state, such that a simplified treatment ignoring any other $|a\rangle$ states yields the correct excitation rate to within around 30%. However, significant contributions also emerge from higher-lying $5sns$ $^3\mathrm{S}_1$ manifolds, as well as from $5snd$ $^3\mathrm{D}_1$, $5p^2$ $^3\mathrm{P}^{\mathrm{e}}_1$ and $5d^2$ $^3\mathrm{P}^{\mathrm{e}}_1$. Due to the sign of the various Wigner $6j$ symbols, not all contributions are in the same direction – some manifolds actually (rather unhelpfully) slow down the three-photon Rabi frequency. In the following derivation, we include contributions from all the intermediate states.

The rest-frame atomic Hamiltonian in the presence of all three probe fields can be written in the dipole approximation as:

$$\hat{H} = \sum_{\psi} E_{\psi} \left|\psi\right\rangle \left\langle\psi\right| - \hat{\mathbf{d}} \cdot \mathbf{E}_1(t) - \hat{\mathbf{d}} \cdot \mathbf{E}_2(t) - \hat{\mathbf{d}} \cdot \mathbf{E}_3(t) \tag{4.44}$$

where the three probing fields can be written in the form:

$$\begin{aligned} \mathbf{E}_1(t) &= \mathbf{E}_1^{(+)} e^{-i\omega_1 t} + \mathbf{E}_1^{(-)} e^{i\omega_1 t} \\ \mathbf{E}_2(t) &= \mathbf{E}_2^{(+)} e^{-i\omega_2 t} + \mathbf{E}_2^{(-)} e^{i\omega_2 t} \\ \mathbf{E}_3(t) &= \mathbf{E}_3^{(+)} e^{-i\omega_3 t} + \mathbf{E}_3^{(-)} e^{i\omega_3 t} \end{aligned} \tag{4.45}$$

When we transform into the laser rotating frame, we find a total Hamiltonian $\hat{H}^{(R)} = \hat{H}_0^{(R)} + \hat{H}_1^{(R)}$, with a bare atomic component given by:

$$\hat{H}_0^{(R)} = -\hbar\Delta \sum_{m} \left|{}^3P_1, m\right\rangle \left\langle{}^3P_1, m\right| - \left(\sum_{a,m} \hbar\Delta_a \left|a, m\right\rangle \left\langle a, m\right| \right) - \hbar\Delta_{\text{clock}} \left|{}^3P_0\right\rangle \left\langle{}^3P_0\right| \tag{4.46}$$

where $\Delta = \omega_1 - \omega_{{}^3P_1}$ is the detuning of the probe field $\mathbf{E}_1$ from the ${}^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_1$ resonance, where $\Delta_a = \omega_1 + \omega_2 - \omega_a$ are the detunings of the sum frequency of fields $\mathbf{E}_1$ and $\mathbf{E}_2$ from the two-photon transition from ${}^1\mathrm{S}_0$ to the auxiliary states $\left|a\right\rangle$, and where $\Delta_{\text{clock}} = \omega_1 + \omega_2 - \omega_3 - \omega_{\text{clock}}$ is the detuning of all three probe beams put together from three-photon clock transition.

Discarding counter-rotating terms, the $E1$-interaction part of the Hamiltonian is given by:

$$\begin{aligned} \hat{H}_1^{(R)} = \Bigg[&- \sum_{m_1} \left\langle{}^1S_0\right| \hat{\mathbf{d}} \cdot \mathbf{E}_1^{(-)} \left|{}^3P_1, m_1\right\rangle \left|{}^1S_0\right\rangle \left\langle{}^3P_1, m_1\right| \\ &- \sum_{a} \sum_{m_1, m_2} \left\langle{}^3P_1, m_1\right| \hat{\mathbf{d}} \cdot \mathbf{E}_2^{(-)} \left|a, m_2\right\rangle \left|{}^3P_1, m_1\right\rangle \left\langle a, m_2\right| \\ &- \sum_{a} \sum_{m_2} \left\langle a, m_2\right| \hat{\mathbf{d}} \cdot \mathbf{E}_3^{(+)} \left|{}^3P_0\right\rangle \left|a, m_2\right\rangle \left\langle{}^3P_0\right| \Bigg] + \text{h.c.} \end{aligned} \tag{4.47}$$

The 689 nm probe field $\mathbf{E}_1$ induces a first-order perturbation to the ground ${}^1\mathrm{S}_0$ state, yielding a dressed state:

$$\left\langle {}^1S_0' \right| = \left\langle {}^1S_0 \right| - \frac{1}{\hbar\Delta} \sum_{m_1} \left\langle {}^1S_0 \right| \hat{\mathbf{d}} \cdot \mathbf{E}_1^{(-)} \left| {}^3P_1, m_1 \right\rangle \left\langle {}^3P_1, m_1 \right| \tag{4.48}$$

Meanwhile, the lattice field $\mathbf{E}_3$ admixes a series of auxiliary states $|a\rangle$ into the ${}^3\mathrm{P}_0$ excited clock state:

$$\left| {}^3P_0' \right\rangle = \left| {}^3P_0 \right\rangle - \sum_{a,m_2} \frac{1}{\hbar(\Delta_a - \Delta_{\mathrm{clock}})} \langle a, m_2 | \, \hat{\mathbf{d}} \cdot \mathbf{E}_3^{(-)} \left| {}^3P_0 \right\rangle |a, m\rangle \tag{4.49}$$

where $\Delta_a - \Delta_{\mathrm{clock}} \approx \Delta_a$ is the detuning of the probe field $\mathbf{E}_3$ from the ${}^3\mathrm{P}_0 \rightarrow |a\rangle$ resonance.

The remaining probe field $\mathbf{E}_2$ then drives a total three-photon Rabi frequency on the clock transition according to:

$$\begin{aligned}
\Omega_{3\gamma} &= -\frac{2 \langle {}^1S_0' | \, \hat{\mathbf{d}} \cdot \mathbf{E}_2^{(-)} \, | {}^3P_0' \rangle}{\hbar} \\
&= -2 \sum_{a,m_1,m_2} \frac{\langle {}^1S_0 | \, \hat{\mathbf{d}} \cdot \mathbf{E}_1^{(-)} \, | {}^3P_1, m_1 \rangle}{\hbar\Delta} \frac{\langle {}^3P_1, m_1 | \, \hat{\mathbf{d}} \cdot \mathbf{E}_2^{(-)} \, | a, m_2 \rangle}{\hbar\Delta_a} \frac{\langle a, m_2 | \, \hat{\mathbf{d}} \cdot \mathbf{E}_3^{(+)} \, | {}^3P_0 \rangle}{\hbar} \\
&= -\frac{2}{\hbar^3 \Delta\Delta_a} \sum_{\substack{a,m_1,m_2 \\ q_1,q_2,q_3}} \Bigg[(-1)^{q_1+q_2+q_3} (\mathbf{E}_1^{(-)})_{-q_1} (\mathbf{E}_2^{(-)})_{-q_2} (\mathbf{E}_3^{(+)})_{-q_3} \\
&\qquad \times \left\langle {}^1S_0 \right| \hat{d}_{q_1} \left| {}^3P_1, m_1 \right\rangle \left\langle {}^3P_1, m_1 \right| \hat{d}_{q_2} \left| a, m_2 \right\rangle \left\langle a, m_2 \right| \hat{d}_{q_3} \left| {}^3P_0 \right\rangle \Bigg]
\end{aligned}$$

which, if we apply the Wigner-Eckart theorem and insert all the Clebsch-Gordan coefficients, gives the final result:

$$\Omega_{3\gamma} = -i \sum_a \frac{2 \langle {}^1S_0 | |\hat{\mathbf{d}}| | {}^3P_1 \rangle \langle {}^3P_1 | |\hat{\mathbf{d}}| | a \rangle \langle {}^3P_0 | |\hat{\mathbf{d}}| | a \rangle}{3\hbar^3 \Delta\Delta_a \sqrt{2}} \mathbf{E}_1^{(-)} \cdot \left(\mathbf{E}_2^{(-)} \times \mathbf{E}_3^{(+)} \right) \tag{4.50}$$

or, if we gather all the reduced matrix elements and detunings into a single coefficient:

$$\frac{\Omega_{3\gamma}}{2\pi} = \zeta \frac{2\pi}{\Delta} \sqrt{I_1 I_2 I_3} \times i \hat{\mathbf{e}}_1 \cdot (\hat{\mathbf{e}}_2 \times \hat{\mathbf{e}}_3^*) \tag{4.51}$$

where ζ is a measure of the induced Rabi frequency per unit of intensity of each probe laser field, and $\hat{\mathbf{e}}_1$, $\hat{\mathbf{e}}_2$, and $\hat{\mathbf{e}}_3$ are the unit complex polarisations of the probe fields. Note that the clock transition cannnont be driven at all if the probe polarisations all lie in the same plane – therefore, we can't have the beams all propagating along the same axis (i.e. at least one beam must come in 'from the side'). In addition, we observe that the Rabi frequency is independent of the direction of the quantisation axis – thus, no bias magnetic field will be required, allowing the straightforward reduction of the quadratic Zeeman shift to a negligible level

Feeding the Sr transition data originally gathered by Ian Hill into equation 4.50 (numerous original sources are listed on page 149 of his thesis [106]), we calculate the three-photon excitation rate coefficient as $\zeta = 27.6\ \text{Hz} \cdot \text{MHz}/(\text{W}\,\text{cm}^{-2})^{3/2}$. Writing this result in a more user-friendly form, and assuming an optimum relative polarisation of the probes such that $\hat{\mathbf{e}}_1 \cdot (\hat{\mathbf{e}}_2 \times \hat{\mathbf{e}}_3^*) = 1$, we find:

$$\Omega_{3\gamma} = \left(\frac{10\,\text{MHz}}{\Delta/2\pi}\right) \times \sqrt{\frac{I_1}{1\,\text{mW/cm}^2} \cdot \frac{I_2}{10\,\text{W/cm}^2} \cdot \frac{I_3}{10\,\text{kW/cm}^2}} \times 27.6\,\text{Hz} \tag{4.52}$$

Since we haven't yet experimentally optimised the three-photon scheme, equation 4.52 indicates our best guess at plausible operating parameters I_1, I_2, I_3, and Δ for a hyper-Ramsey ^{88}Sr clock. With these intensities we expect to drive a Rabi $\pi/2$-pulse in around 9 ms, even considering the fact that we've chosen a fairly low lattice intensity $I_3 = 10\,\text{kHz}$ corresponding to a modest trap depth $U_0 \approx 40E_r$. The two largest probe-related systematics are the Stark shifts from the 689 nm and 826 nm probe beams (fields $\mathbf{E}_1$ and $\mathbf{E}_2$ respectively), which are governed by the equations:

$$\Delta\nu_{\mathbf{E}_1} = -\left(\frac{10\,\text{MHz}}{\Delta} \cdot \frac{I_1}{1\,\text{mW/cm}^2}\right) \times 231\,\text{Hz} \tag{4.53}$$

$$\Delta\nu_{\mathbf{E}_2} = \left(\frac{I_1}{10\,\text{W/cm}^2}\right) \times 4.7\,\text{Hz} \tag{4.54}$$

both of which could be adequately suppressed using hyper-Ramsey spectroscopy. If necessary, there is potentially room for significantly reducing $\Delta_{\mathbf{E}_1}$, either by extending the pulse time or by using a smaller detuning Δ from the cooling transition – these parameters will have to be optimised in experiment.

There are also a few additional nasty effects caused by the admixture of $^3\text{P}_1$ into the clock states. The size of the admixtures is governed by the relations:

$$\left|\left\langle {}^1S_0'\middle|\, {}^3P_1\right\rangle\right|^2 = \left(\frac{\Omega_{^1S_0\to{}^3P_1}}{2\Delta}\right)^2 = \left(\frac{10\,\mathrm{MHz}}{\Delta}\right)^2\left(\frac{I_1}{1\,\mathrm{mW/cm^2}}\right)\times 2.3\times 10^{-5} \tag{4.55}$$

$$\left|\left\langle {}^3P_0'\middle|\, {}^3P_1\right\rangle\right|^2 = \left(\frac{\Omega_{^3P_0\to{}^3P_1}}{2\Delta}\right)^2 = \left(\frac{10\,\mathrm{MHz}}{\Delta}\right)^2\left(\frac{I_2}{10\,\mathrm{W/cm^2}}\cdot\frac{I_3}{10\,\mathrm{kW/cm^2}}\right)\times 8\times 10^{-8} \tag{4.56}$$

where the small admixture of $^3\mathrm{P}_1$ into $^3\mathrm{P}_0'$ is due to a second-order perturbation in $|\mathbf{E}_2|\,|\mathbf{E}_3|$, with $\Omega^2_{^3P_0\to{}^3P_1}$ denoting a two-photon Rabi frequency. We have assumed that $|\hat{\mathbf{e}}_2\times\hat{\mathbf{e}}_3^*| = 1$, and we have summed over the Zeeman sublevels $m = -1, 0, 1$. Clearly, for these operating conditions (and for most realistic scenarios), the dominant admixture is caused by the 689 nm probe beam.

The state admixtures have two important parasitic effects on the three-photon scheme. Firstly, photons are scattered on the $^3\mathrm{P}_1 \to {}^1\mathrm{S}_0$ transition, causing atomic heating and decoherence. The scatter rate from the dressed ground state is given by $\Gamma_{\mathrm{sc}} = |\langle^1S_0'|\,^3P_1\rangle|^2\times\Gamma_{^3P_1} = |\langle^1S_0'|\,^3P_1\rangle|^2\times 4.8\times10^4\,\mathrm{s}^{-1}$, which gives us $\Gamma_{\mathrm{sc}}\approx 1.1\,\mathrm{s}^{-1}$ for our default operating conditions. This ought to have very little effect within the short 9 ms hyper-Ramsey $\pi/2$-pulses.

The second problem is the hyperpolarisability shift induced by the 689 nm and 813 nm fields, scaling as $|\mathbf{E}_1|^2\,|\mathbf{E}_3|^2$, with an associated broadening due to the spatial distribution of the lattice intensity I_3 within the lattice trap. This shift will general be dependent on the polarisations $\hat{\mathbf{e}}_1$ and $\hat{\mathbf{e}}_3$, but we crudely estimate the worst case here: assuming the least-magic polarisation, the differential shift at 813 nm to the $^1\mathrm{S}_0\to{}^3\mathrm{P}_1$ transition is approximately half the depth of the lattice [43], giving a differential shift of 50 kHz for $I_3 = 10\,\mathrm{kW\,cm^{-2}}$. This ought to translate into a hyperpolarisability shift to the clock transition below $2.3\times10^{-5}\times 50\,\mathrm{kHz}\approx 1\,\mathrm{Hz}$ – sizeable, but routinely suppressed by the hyper-Ramsey technique.

To summarise, we have found that the shifts to the clock transition are all at a manageable level in the three-photon excitation scheme, and that they are all amenable to suppression by hyper-Ramsey spectroscopy. In addition, we note that the presence of three photons has a nice bonus: it enables the alignment of the probe beams into a Doppler-free orientation where the net photon recoil $\hbar\mathbf{k}_1+\hbar\mathbf{k}_2-\hbar\mathbf{k}_3 = 0$. The full cancellation of recoil ensures that there is no coupling of the clock excitation to the atomic motional degree of freedom, meaning that the atoms no longer need to be as tightly confined in order to reach the Lamb-Dicke regime. In particular, a Doppler-free lattice clock would be able to use much lower lattice depths down to

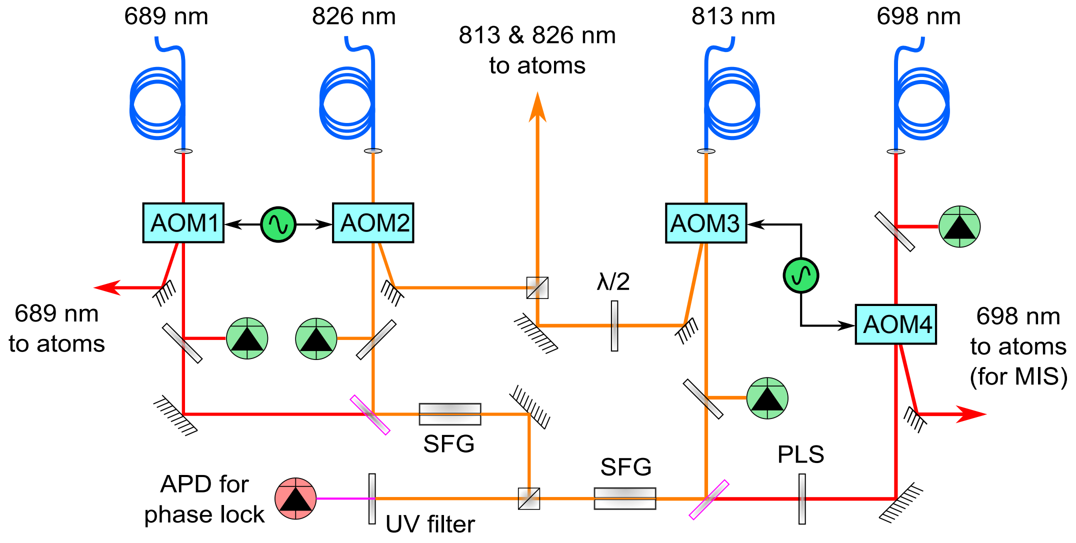

Figure 4.12: Laser system for 3 photon spectroscopy of ^{88}Sr. SFG: Sum frequency generation crystal. PLS: Sampler for path length stabilisation, MIS: Magnetically-induced spectroscopy. Amplified Si photodiodes for intensity stabilisation are shown in green. Additional AOMs (not shown) for intensity stabilisation are placed before each fibre.

$\sim 5E_r$ without needing to worry about frequency shifts from site-to-site tunnelling[1] [153].

Now we turn to the problem of experimental realisation of the three-photon scheme. Naturally we already have the required lasers at 689 nm and 813 nm, ordinarily used for 2$^{\text{nd}}$-stage cooling and lattice trapping respectively. Therefore, only one additional laser was needed, at the not-too-difficult wavelength of 826 nm – thankfully, we had a nice old lattice ECDL lying in the cupboard, historically made redundant by our Ti:Sapphire laser (see section 2.3.3.4), into which we have now installed a Thorlabs L830P200 200 mW diode. The diode's natural lasing wavelength turned out to be 823 nm, so we had no trouble pulling it to 826 nm using the grating angle.

The only major challenge remaining is the frequency stabilisation of the three probe beams such that $\omega_1 + \omega_2 - \omega_3 = \omega_{\text{clock}}$. The required specifications are quite exacting: we need to make sure that the sum frequency maintains a coherence time comparable to our present local oscillator at 698 nm, so that we can drive three-photon excitation pulses lasting 100 ms or more. One option would be to stabilise each of the

---

[1] A slightly higher lattice depth than $5E_r$ may still be needed in order to retain a reasonable lifetime for the Wannier-Stark states, and to avoid significant collisional shifts between atoms in neighbouring lattice sites.

three probe lasers independently at the Hz level; however, building three ultrastable local oscillators at different wavelengths would be a pretty extreme task.

Instead, we plan to implement the phase-locking scheme depicted in figure 4.12. We will generate two sum frequencies $\omega_1 + \omega_2$ and $\omega_3 + \omega_{clock}$ using two custom waveguide ppKTP chips from AdvR, resulting in two output beams at 375.8 nm. Extracting the beat frequency on an APD, we can phase lock Δ_{beat} to a local RF reference by actuating on the 826 nm laser – the net result will impose a fixed relation $\omega_1 + \omega_2 = \omega_3 + \omega_{\text{clock}} + \Delta_{\text{beat}}$. The beams sent into the phase-locking scheme are deliberately chosen as the zeroth orders passing through the final probe 'switching' AOMs – that way, any thermal AOM-induced phase chirps should be 'seen' by Δ_{beat} just the same as by the atoms, so they should get largely removed by the phase lock (see section 4.1.6). The same zero-order beams are also sampled onto continuously-operated intensity-stabilisation photodiodes, preventing significant perturbations to the phase lock when the probe light is switched on and off.

Finally, we ensure that the optical setup is still capable of implementing the more conventional method of magnetically-induced spectroscopy. Therefore, we should be able to interleave consecutive sequences of (1) three-photon and (2) magnetically-induced spectroscopy in order to evaluate any frequency offsets between the two schemes – since each scheme has an entirely different set of probe-induced shifts, this should give an independent verification of the accuracy of both techniques. Hopefully they will agree!

4.3 The future of lattice clocks

The recent pace of progress in optical atomic clocks has been incredibly fast. Before the year 2000, when the first octave-spanning frequency combs were developed, the most precisely-known optical frequencies still had uncertainties in the 10^{-13} range [24, 238]. Since then, the inaccuracy of the best optical clocks has been plummeting roughly by a factor of ten every 4 years (see chapter 1), culminating in JILA's 2015 declaration of an Sr lattice clock with a total systematic uncertainty of just 2×10^{-18} [193][1].

For better or worse, there has been less fanfare about the other exceptional feature of optical atomic clocks – their stability. In many applications, however, a clock's stability is even more important than its accuracy: for instance, if we need a precise

---

[1] The goalposts have moved especially quickly for optical lattice clocks: when I first started my research in late 2012 (3.5 years ago), the best reported inaccuracy of a lattice clock was still at 1.5×10^{-16} [167].

measurement *resolution* within a given time τ, then the clock's instability $\sigma_y(\tau)$ is the important limiting factor. Once again, the most recent generation of lattice clocks now give the best performance of any clock species, with instabilities on the order of $\sigma_y(\tau) \approx 2 \times 10^{-16}/\sqrt{\tau}$ [6, 193].

Where will this progress end? In this chapter, we have discovered that no fundamental barriers exist to prevent further reductions in the systematic uncertainty of optical lattice clocks: with some careful engineering of the black-body environment, and with the use of shallow, well-controlled lattice traps, it looks perfectly achievable to reach total uncertainties well below 1×10^{-18}. Earlier, in chapter 3, we also described a plausible route towards improvements in clock stability: a composite lattice clock with zero dead time could eliminate the nefarious Dick effect, potentially enabling quantum noise-limited instability of $\sigma_y(\tau) \approx 3 \times 10^{-17}/\sqrt{\tau}$ for typical lattice clocks using 1000 atoms and 300 ms pulses. If the Dick-free lattice clock were upgraded to use a few 10^5 atoms in a 3D lattice, or to use 10^3 atoms prepared into a spin-squeezed state, then it could even approach incredible quantum-limited instabilities of $\sigma_y(\tau) \lesssim 3 \times 10^{-18}/\sqrt{\tau}$.

The realisation of such an advanced optical lattice clock, capable of 19-digit precision within a few minutes of measurement, looks almost inevitable over the next few years. What could we do with clocks that can routinely explore Earth's gravitational field to a height precision of less than 1 mm? I don't know, but it will be interesting to find out.

Appendix A

My contribution to the project

When I started on the lattice clock project at NPL, we hadn't yet realised a red MOT. In fact, the old MOT coils had just been removed from the main chamber, so we couldn't make blue MOTs either. Since then, I have helped to transform the experiment into a fully fledged, highly reliable ^{88}Sr/^{87}Sr optical lattice clock.

Various milestones were reached while I was in the lab running the experiment. In chronological order: the first blue MOT using the new MOT coils; the first ^{88}Sr red OT (a mysterious blob); the first proper ^{88}Sr red MOT (a few months later!); the first lattice-trapped atoms using the old 813 nm TA (just before it died); the first ^{87}Sr red MOT; the first sighting of the clock transition (in the ^{87}Sr red MOT); the first lattice-trapped atoms using the new 813 nm Ti:sapphire laser; the first clock spectroscopy in the lattice; the first 'streak loading' of ^{88}Sr into the lattice; the first lock of the clock laser to the atomic transition; the first major frequency measurement campaign (ITOC, May 2015); the first systematic evaluation; the first use of the universal synthesiser to reduce the clock instability; the first results from hyper-Ramsey spectroscopy.

Of course, the vast majority of the labour that goes into a project like this isn't actually running the experiment; instead, most of the time is spent building the components that make it work. I give a rather incomplete list of my contributions here. My first job was to design and build the flexible driving circuitry for the new high-current MOT coils (see section 2.2.4.1). Later, I applied this electronics experience to design, build and test the fast/general-purpose loop filters described in section 2.2.3. I installed and tuned most of the 30 such devices now in use, along with the various associated bits of optics and opto-electronics which they control. I wrote the front-end labscript program now used to run the experiment (see section 2.2.2.2), and resolved various bugs that initially plagued the compatibility of the labscript suite with the FPGA hardware. I also wrote various python programs to process, plot, and fit to the experimental data. I extended the 689 nm optics system, including the construction of slave lasers for cooling and spin-polarising ^{87}Sr, and including the lock to the transfer cavity (RAM servo and all, see section 2.2.1.2). I designed and built the 'Sr lattice lab' side of the universal synthesiser feedforward scheme in section 3.1.3. I also had to build, extend, or upgrade various other laser and optics systems over the years to meet our ever-evolving requirements.

Coming up with atomic physics ideas is also an important part of the job, which, perhaps unfairly, gets rewarded with the highest-impact publications. In this regard, I invented, modelled, and implemented the modified hyper-Ramsey spectroscopy scheme described in 4.2.2, and I wrote the paper on the topic which was ultimately published as a PRA rapid communication [110]. I also invented, modelled, and began building the apparatus for the multi-photon excitation schemes in section 4.2.3. Finally, I wrote the python modelling program behind all the Dick effect results in section 3.3, enabling the characterisation of the instability of our optical lattice clock.

Bibliography

[1] D N Aguilera, H Ahlers, B Battelier, A Bawamia, A Bertoldi, R Bondarescu, K Bongs, P Bouyer, C Braxmaier, L Cacciapuoti, C Chaloner, M Chwalla, W Ertmer, M Franz, N Gaaloul, M Gehler, D Gerardi, L Gesa, N Grlebeck, J Hartwig, M Hauth, O Hellmig, W Herr, S Herrmann, A Heske, A Hinton, P Ireland, P Jetzer, U Johann, M Krutzik, A Kubelka, C Lmmerzahl, A Landragin, I Lloro, D Massonnet, I Mateos, A Milke, M Nofrarias, M Oswald, A Peters, K Posso-Trujillo, E Rasel, E Rocco, A Roura, J Rudolph, W Schleich, C Schubert, T Schuldt, S Seidel, K Sengstock, C F Sopuerta, F Sorrentino, D Summers, G M Tino, C Trenkel, N Uzunoglu, W von Klitzing, R Walser, T Wendrich, A Wenzlawski, P Weels, A Wicht, E Wille, M Williams, P Windpassinger, and N Zahzam. STE-QUEST - test of the universality of free fall using cold atom interferometry. *Classical and Quantum Gravity*, 31(11):115010, 2014.

[2] Daisuke Akamatsu, Hajime Inaba, Kazumoto Hosaka, Masami Yasuda, Atsushi Onae, Tomonari Suzuyama, Masaki Amemiya, and Feng-Lei Hong. Spectroscopy and frequency measurement of the ^{87}Sr clock transition by laser linewidth transfer using an optical frequency comb. *Applied Physics Express*, 7(1):012401, 2014.

[3] Daisuke Akamatsu, Masami Yasuda, Hajime Inaba, Kazumoto Hosaka, Takehiko Tanabe, Atsushi Onae, and Feng-Lei Hong. Frequency ratio measurement of ^{171}Yb and ^{87}Sr optical lattice clocks. *Opt. Express*, 22(7):7898–7905, Apr 2014.

[4] T. Akatsuka, M. Takamoto, and H. Katori. Optical lattice clocks with non-interacting bosons and fermions. *Nature Physics*, 4(12):954–959, December 2008.

[5] Tomoya Akatsuka, Masao Takamoto, and Hidetoshi Katori. Three-dimensional

optical lattice clock with bosonic ^{88}Sr atoms. *Phys. Rev. A*, 81:023402, Feb 2010.

[6] Ali Al-Masoudi, Sören Dörscher, Sebastian Häfner, Uwe Sterr, and Christian Lisdat. Noise and instability of an optical lattice clock. *Phys. Rev. A*, 92:063814, Dec 2015.

[7] E. A. Alden, K. R. Moore, and A. E. Leanhardt. Two-photon $E1$-$M1$ optical clock. *Phys. Rev. A*, 90:012523, Jul 2014.

[8] J. Alnis, A. Matveev, N. Kolachevsky, Th. Udem, and T.W. Hänsch. Subhertz linewidth diode lasers by stabilization to vibrationally and thermally compensated ultralow-expansion glass Fabry-Pérot cavities. *Physical Review A*, 77:053809, May 2008.

[9] Brett Altschul, Quentin G. Bailey, Luc Blanchet, Kai Bongs, Philippe Bouyer, Luigi Cacciapuoti, Salvatore Capozziello, Naceur Gaaloul, Domenico Giulini, Jonas Hartwig, Luciano Iess, Philippe Jetzer, Arnaud Landragin, Ernst Rasel, Serge Reynaud, Stephan Schiller, Christian Schubert, Fiodor Sorrentino, Uwe Sterr, Jay D. Tasson, Guglielmo M. Tino, Philip Tuckey, and Peter Wolf. Quantum tests of the einstein equivalence principle with the stequest space mission. *Advances in Space Research*, 55(1):501 – 524, 2015.

[10] John D. Anderson, Philip A. Laing, Eunice L. Lau, Anthony S. Liu, Michael Martin Nieto, and Slava G. Turyshev. Indication, from pioneer 10/11, galileo, and ulysses data, of an apparent anomalous, weak, long-range acceleration. *Phys. Rev. Lett.*, 81:2858–2861, Oct 1998.

[11] M.H. Anderson, J.R. Ensher, M.R. Matthews, C.E. Wieman, and E.A. Cornell. Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor. *Science*, 269(5221):198–201, 1995.

[12] E. J. Angstmann, V. A. Dzuba, and V. V. Flambaum. Relativistic effects in two valence-electron atoms and ions and the search for variation of the fine-structure constant. *Phys. Rev. A*, 70:014102, Jul 2004.

[13] Julija Bagdonaite, Paul Jansen, Christian Henkel, Hendrick L. Bethlem, Karl M. Menten, and Wim Ubachs. A stringent limit on a drifting proton-to-electron mass ratio from alcohol in the early universe. *Science*, 339(6115):46–48, 2013.

[14] X. Baillard, A. Gauguet, S. Bize, P. Lemonde, Ph. Laurent, A. Clairon, and P. Rosenbusch. Interference-filter-stabilized external-cavity diode lasers. *Optics Communications*, 266(2):609 – 613, 2006.

[15] Z. W. Barber, C. W. Hoyt, C. W. Oates, L. Hollberg, A. V. Taichenachev, and V. I. Yudin. Direct excitation of the forbidden clock transition in neutral ^{174}Yb atoms confined to an optical lattice. *Phys. Rev. Lett.*, 96:083002, Mar 2006.

[16] Z. W. Barber, J. E. Stalnaker, N. D. Lemke, N. Poli, C. W. Oates, T. M. Fortier, S. A. Diddams, L. Hollberg, C. W. Hoyt, A. V. Taichenachev, and V. I. Yudin. Optical Lattice Induced Light Shifts in an Yb Atomic Clock. *Phys. Rev. Lett.*, 100:103002, Mar 2008.

[17] Z.W. Barber. *Ytterbium Optical Lattice Clock*. PhD thesis, University of Colorado, 2007.

[18] G. P. Barwood, G. Huang, H. A. Klein, L. A. M. Johnson, S. A. King, H. S. Margolis, K. Szymaniec, and P. Gill. Agreement between two ^{88}Sr^{+} optical clocks to 4 parts in 10^{17}. *Phys. Rev. A*, 89:050501, May 2014.

[19] K. Beloy, N. Hinkley, N. B. Phillips, J. A. Sherman, M. Schioppo, J. Lehman, A. Feldman, L. M. Hanssen, C. W. Oates, and A. D. Ludlow. Atomic clock with 1×10^{-18} room-temperature blackbody stark uncertainty. *Phys. Rev. Lett.*, 113:260801, Dec 2014.

[20] K. Beloy, M. G. Kozlov, A. Borschevsky, A. W. Hauser, V. V. Flambaum, and P. Schwerdtfeger. Rotational spectrum of the molecular ion NH^{+} as a probe for α and m_e/m_p variation. *Phys. Rev. A*, 83:062514, Jun 2011.

[21] K. Beloy, J. A. Sherman, N. D. Lemke, N. Hinkley, C. W. Oates, and A. D. Ludlow. Determination of the $5d6s$ $^3D_1$ state lifetime and blackbody-radiation clock shift in yb. *Phys. Rev. A*, 86:051404, Nov 2012.

[22] J. C. Berengut, V. A. Dzuba, and V. V. Flambaum. Enhanced laboratory sensitivity to variation of the fine-structure constant using highly charged ions. *Phys. Rev. Lett.*, 105:120801, Sep 2010.

[23] J.C. Bergquist, W.M. Itano, and D.J. Wineland. Laser Stabilization to a Single Ion. In *Proceedings of the International School of Physics, Enrico Fermi,Italy, June 23-July 3, 1992*, 1996.

[24] J. E. Bernard, A. A. Madej, L. Marmet, B. G. Whitford, K. J. Siemsen, and S. Cundy. Cs-based frequency measurement of a single, trapped ion transition in the visible region of the spectrum. *Phys. Rev. Lett.*, 82:3228–3231, Apr 1999.

[25] J. W. Berthold and S. F. Jacobs. Ultraprecise thermal expansion measurements of seven low expansion materials. *Applied Optics*, 15(10):2344–2347, 1976.

[26] G. W. Biedermann, K. Takase, X. Wu, L. Deslauriers, S. Roy, and M. A. Kasevich. Zero-dead-time operation of interleaved atomic clocks. *Phys. Rev. Lett.*, 111:170802, Oct 2013.

[27] T. Binnewies, G. Wilpers, U. Sterr, F. Riehle, J. Helmcke, T. E. Mehlstäubler, E. M. Rasel, and W. Ertmer. Doppler cooling and trapping on forbidden transitions. *Phys. Rev. Lett.*, 87:123002, Aug 2001.

[28] M. Bishof, Y. Lin, M. D. Swallows, A. V. Gorshkov, J. Ye, and A. M. Rey. Resolved Atomic Interaction Sidebands in an Optical Clock Transition. *Physical Review Letters*, 106:250801, Jun 2011.

[29] M. Bishof, M. J. Martin, M. D. Swallows, C. Benko, Y. Lin, G. Quéméner, A. M. Rey, and J. Ye. Inelastic collisions and density-dependent excitation suppression in a ^{87}Sr optical lattice clock. *Phys. Rev. A*, 84:052716, Nov 2011.

[30] Eric D. Black. An introduction to PoundDreverHall laser frequency stabilization. *American Journal of Physics*, 69(1):79–87, 2001.

[31] S. Blatt, A. D. Ludlow, G. K. Campbell, J. W. Thomsen, T. Zelevinsky, M. M. Boyd, J. Ye, X. Baillard, M. Fouché, R. Le Targat, A. Brusch, P. Lemonde, M. Takamoto, F.-L. Hong, H. Katori, and V. V. Flambaum. New limits on coupling of fundamental constants to gravity using ^{87}Sr optical lattice clocks. *Phys. Rev. Lett.*, 100:140801, Apr 2008.

[32] S. Blatt, J. W. Thomsen, G. K. Campbell, A. D. Ludlow, M. D. Swallows, M. J. Martin, M. M. Boyd, and J. Ye. Rabi spectroscopy and excitation inhomogeneity in a one-dimensional optical lattice clock. *Phys. Rev. A*, 80:052703, Nov 2009.

[33] B. J. Bloom, T. L. Nicholson, J. R. Williams, S. L. Campbell, M. Bishof, X. Zhang, W. Zhang, S. L. Bromley, and J. Ye. An optical lattice clock with accuracy and stability at the 10^{-18} level. *Nature*, 506(7486):71–75, 02 2014.

[34] Marcin Bober, Piotr Morzyski, Agata Cygan, Daniel Lisak, Piotr Masowski, Mateusz Prymaczek, Piotr Wciso, Piotr Ablewski, Mariusz Piwiski, Szymon Wjtewicz, Katarzyna Bielska, Dobrosawa Bartoszek-Bober, Ryszard S Trawiski, Micha Zawada, Roman Ciuryo, Jerzy Zachorowski, Marcin Piotrowski, Wojciech Gawlik, Filip Ozimek, and Czesaw Radzewicz. Strontium optical lattice clocks for practical realization of the metre and secondary representation of the second. *Measurement Science and Technology*, 26(7):075201, 2015.

[35] J.G. Bohnet. *A Superradiant Laser and Spin Squeezed States: Collective Phenomena in a Rubidium Cavity QED System for Enhancing Precision Measurements*. PhD thesis, University of Colorado, 2014.

[36] Justin G. Bohnet, Zilong Chen, Joshua M. Weiner, Dominic Meiser, Murray J. Holland, and James K. Thompson. A steady-state superradiant laser with less than one intracavity photon. *Nature*, 484(7392):78–81, Apr 2012.

[37] Ruxandra Bondarescu, Mihai Bondarescu, Gyrgy Hetnyi, Lapo Boschi, Philippe Jetzer, and Jayashree Balakrishna. Geophysical applicability of atomic clocks: direct continental geoid mapping. *Geophysical Journal International*, 191(1):78–82, 2012.

[38] Ruxandra Bondarescu, Andreas Schrer, Andrew Lundgren, Gyrgy Hetnyi, Nicolas Houli, Philippe Jetzer, and Mihai Bondarescu. Ground-based optical atomic clocks as a tool to monitor vertical surface motion. *Geophysical Journal International*, 202(3):1770–1774, 2015.

[39] Kai Bongs, Yeshpal Singh, Lyndsie Smith, Wei He, Ole Kock, Dariusz wierad, Joshua Hughes, Stephan Schiller, Soroosh Alighanbari, Stefano Origlia, Stefan Vogt, Uwe Sterr, Christian Lisdat, Rodolphe Le Targat, Jrme Lodewyck, David Holleville, Bertrand Venon, Sbastien Bize, Geoffrey P. Barwood, Patrick Gill, Ian R. Hill, Yuri B. Ovchinnikov, Nicola Poli, Guglielmo M. Tino, Jrgen Stuhler, and Wilhelm Kaenders. Development of a strontium optical lattice clock for the {SOC} mission on the {ISS}. *Comptes Rendus Physique*, 16(5):553 – 564, 2015. The measurement of time / La mesure du temps.

[40] J. Borregaard and A. S. Sørensen. Efficient atomic clocks operated with several atomic ensembles. *Phys. Rev. Lett.*, 111:090802, Aug 2013.

[41] Benjamin Botermann, Dennis Bing, Christopher Geppert, Gerald Gwinner, Theodor W. Hänsch, Gerhard Huber, Sergei Karpuk, Andreas Krieger, Thomas

Kühl, Wilfried Nörtershäuser, Christian Novotny, Sascha Reinhardt, Rodolfo Sánchez, Dirk Schwalm, Thomas Stöhlker, Andreas Wolf, and Guido Saathoff. Test of Time Dilation Using Stored Li^+ Ions as Clocks at Relativistic Speed. *Phys. Rev. Lett.*, 113:120405, Sep 2014.

[42] Martin M. Boyd, Andrew D. Ludlow, Sebastian Blatt, Seth M. Foreman, Tetsuya Ido, Tanya Zelevinsky, and Jun Ye. $^{87}\mathrm{Sr}$ lattice clock with inaccuracy below 10^{-15}. *Phys. Rev. Lett.*, 98:083002, Feb 2007.

[43] M.M. Boyd. *High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard for Frequency and Time.* PhD thesis, University of Colorado, 2007.

[44] M.M. Boyd, T. Zelevinsky, A.D. Ludlow, S. Blatt, T. Zanon-Willette, S.M. Foreman, and Jun Ye. Nuclear spin effects in optical lattice clocks. *Physical Review A*, 76(2):022510, Aug 2007.

[45] K.-A. Brickman, M.-S. Chang, M. Acton, A. Chew, D. Matsukevich, P.C. Haljan, V.S. Bagnato, and C. Monroe. Magneto-optical trapping of cadmium. *Physical Review A*, 76(4):043411, Oct 2007.

[46] E.M. Bridge. *Towards a Strontium Optical Lattice Clock.* PhD thesis, University of Oxford, 2013.

[47] J.R. Brophy, L. Friedman, and F. Culick. Asteroid retrieval feasibility. In *Aerospace Conference, 2012 IEEE*, pages 1–16, March 2012.

[48] A. Brusch, R. Le Targat, X. Baillard, M. Fouche, and P. Lemonde. Hyperpolarizability Effects in a Sr Optical Lattice Clock. *Physical Review Letters*, 96(10):103003, 2006.

[49] L. Cacciapuoti, N. Dimarcq, G. Santarelli, P. Laurent, P. Lemonde, A. Clairon, P. Berthoud, A. Jornod, F. Reina, S. Feltham, and C. Salomon. Atomic clock ensemble in space: Scientific objectives and mission status. *Nuclear Physics B - Proceedings Supplements*, 166:303 – 306, 2007. Proceedings of the Third International Conference on Particle and Fundamental Physics in SpaceProceedings of the Third International Conference on Particle and Fundamental Physics in Space.

[50] G. K. Campbell, M. M. Boyd, J. W. Thomsen, M. J. Martin, S. Blatt, M. D. Swallows, T. L. Nicholson, T. Fortier, C. W. Oates, S. A. Diddams, N. D.

Lemke, P. Naidon, P. Julienne, Jun Ye, and A. D. Ludlow. Probing interactions between ultracold fermions. *Science*, 324(5925):360–363, 2009.

[51] Y. Castin, H. Wallis, and J. Dalibard. Limit of Doppler cooling. *Journal of the Optical Society of America B*, 6(11):2046–2057, Nov 1989.

[52] Lisheng Chen, John L. Hall, Jun Ye, Tao Yang, Erjun Zang, and Tianchu Li. Vibration-induced elastic deformation of fabry-perot cavities. *Phys. Rev. A*, 74:053801, Nov 2006.

[53] Cheng Chin, V V Flambaum, and M G Kozlov. Ultracold molecules: new probes on the variation of fundamental constants. *New Journal of Physics*, 11(5):055048, 2009.

[54] C. W. Chou, D. B. Hume, J. C. J. Koelemeij, D. J. Wineland, and T. Rosenband. Frequency Comparison of Two High-Accuracy Al^+ Optical Clocks. *Phys. Rev. Lett.*, 104:070802, Feb 2010.

[55] C. W. Chou, D. B. Hume, M. J. Thorpe, D. J. Wineland, and T. Rosenband. Quantum coherence between two atoms beyond $q = 10^{15}$. *Phys. Rev. Lett.*, 106:160801, Apr 2011.

[56] C.W. Chou, D.B. Hume, T. Rosenband, and D.J. Wineland. Optical Clocks and Relativity. *Science*, 329(5999):1630–1633, 2010.

[57] M. Chwalla, J. Benhelm, K. Kim, G. Kirchmair, T. Monz, M. Riebe, P. Schindler, A.S. Villar, W. Hänsel, C.F. Roos, R. Blatt, M. Abgrall, G. Santarelli, G.D. Rovera, and Ph. Laurent. Absolute Frequency Measurement of the $^{40}\mathrm{Ca}^+$ 4s $^2\mathrm{S}_{1/2}$ - 3d $^2\mathrm{D}_{5/2}$ Clock Transition. *Physical Review Letters*, 102(2):023002, Jan 2009.

[58] M. Chwalla, K. Kim, T. Monz, P. Schindler, M. Riebe, C.F. Roos, and R. Blatt. Precision spectroscopy with two correlated atoms. *Applied Physics B*, 89(4):483–488, 2007.

[59] C. Cohen-Tannoudji. *Quantum Mechanics Vol. 1.* Wiley, 1977.

[60] G. Cole. Cavity optomechanics with low-noise crytalline mirrors. *Proceedings of SPIE*, 8458(845807), 2012.

[61] G. D. Cole, W. Zhang, M. J. Martin, J. Ye, and M. Aspelmayer. Tenfold reduction of brownian noise in high-reflectivity optical coatings. *Nature Photonics*, 7:644–650, August 2013.

[62] Alexander D. Cronin, Jörg Schmiedmayer, and David E. Pritchard. Optics and interferometry with atoms and molecules. *Rev. Mod. Phys.*, 81:1051–1129, Jul 2009.

[63] E.A. Curtis, C.W. Oates, and L. Hollberg. Quenched narrow-line second- and third-stage laser cooling of ^{40}Ca. *Journal of the Optical Society of America B*, 20(5):977, 2003.

[64] S. De, U. Dammalapati, K. Jungmann, and L. Willmann. Magneto-optical trapping of barium. *Phys. Rev. A*, 79:041402, Apr 2009.

[65] C. Degenhardt, T. Nazarova, C. Lisdat, H. Stoehr, U. Sterr, and F. Riehle. Influence of chirped excitation pulses in an optical clock with ultracold calcium atoms. *IEEE Transactions on Instrumentation and Measurement*, 54(2):771–775, April 2005.

[66] H.G. Dehmelt. Monoion oscillator as potential ultimate laser frequency standard. *Instrumentation and Measurement, IEEE Transactions on*, IM-31(2):83–87, June 1982.

[67] A. Derevianko, B. Obreshkov, and V. A. Dzuba. Mapping out atom-wall interaction with atomic clocks. *Phys. Rev. Lett.*, 103:133201, Sep 2009.

[68] A. Derevianko and M. Pospelov. Hunting for topological dark matter with atomic clocks. *Nat Phys*, 10(12):933–936, Dec 2014. Letter.

[69] G. Dick. Local oscillator induced instabilities in trapped ion frequency standards. In *Precise Time and Time Interval*, pages 133–147, Redondo Beach, 1987.

[70] R. H. Dicke. Coherence in spontaneous radiation processes. *Physical Review*, 93(1):99–110, January 1954.

[71] Scott A. Diddams, David J. Jones, Jun Ye, Steven T. Cundiff, John L. Hall, Jinendra K. Ranka, Robert S. Windeler, Ronald Holzwarth, Thomas Udem, and T. W. Hänsch. Direct Link between Microwave and Optical Frequencies

with a 300 THz Femtosecond Laser Comb. *Phys. Rev. Lett.*, 84:5102–5105, May 2000.

[72] T. H. Dinh, A. Dunning, V. A. Dzuba, and V. V. Flambaum. Sensitivity of hyperfine structure to nuclear radius and quark mass variation. *Phys. Rev. A*, 79:054102, May 2009.

[73] R.W.P. Drever, J.L. Hall, F.V. Kowalski, J. Hough, G.M. Ford, A.J. Munley, and H. Ward. Laser phase and frequency stabilization using an optical resonator. *Applied Physics B: Lasers and Optics*, 31:97–105, 1983. 10.1007/BF00702605.

[74] V. A. Dzuba and V. V. Flambaum. Relativistic corrections to transition frequencies of Ag *i*, Dy *i*, Ho *i*, Yb *ii*, Yb *iii*, Au *i*, and Hg *ii* and search for variation of the fine-structure constant. *Phys. Rev. A*, 77:012515, Jan 2008.

[75] V. A. Dzuba, V. V. Flambaum, and Hidetoshi Katori. Optical clock sensitive to variations of the fine-structure constant based on the Ho^{14+} ion. *Phys. Rev. A*, 91:022119, Feb 2015.

[76] V. A. Dzuba, V. V. Flambaum, M. S. Safronova, S. G. Porsev, T. Pruttivarasin, M. A. Hohensee, and H. Hffner. Strongly enhanced effects of Lorentz symmetry violation in entangled Yb^+ ions. *ArXiv e-prints*, July 2015.

[77] W. Ertmer, C. Schubert, T. Wendrich, M. Gilowski, M. Zaiser, T.v. Zoest, E. Rasel, Ch.J. Bord, A. Clairon, Landragin, P. Laurent, P. Lemonde, G. Santarelli, W. Schleich, F.S. Cataliotti, M. Inguscio, N. Poli, F. Sorrentino, C. Modugno, G.M. Tino, P. Gill, H. Klein, H. Margolis, S. Reynaud, C. Salomon, A. Lambrecht, E. Peik, C. Jentsch, U. Johann, A. Rathke, P. Bouyer, L. Cacciapuoti, P. De Natale, B. Christophe, B. Foulon, P. Touboul, L. Maleki, N. Yu, S.G. Turyshev, J.D. Anderson, F. Schmidt-Kaler, R. Walser, J. Vigu, M. Bchner, M.-C. Angonin, P. Delva, P. Tourrenc, R. Bingham, B. Kent, A. Wicht, L.J. Wang, K. Bongs, Hj. Dittus, C. Lmmerzahl, S. Theil, K. Sengstock, A. Peters, T. Mller, M. Arndt, L. Iess, F. Bondu, A. Brillet, E. Samain, M.L. Chiofalo, F. Levi, and D. Calonico. Matter wave explorer of gravity (mwxg). *Experimental Astronomy*, 23(2):611–649, 2009.

[78] L. Essen and J.V.L. Parry. An Atomic Standard of Frequency and Time Interval: A Caesium Resonator. *Nature*, 176(4476):280–282, August 1955.

[79] M. Evans, S. Ballmer, M. Fejer, P. Fritschel, G. Harry, and G. Ogin. Thermo-optic noise in coated mirrors for high-precision optical measurements. *Phys. Rev. D*, 78:102003, Nov 2008.

[80] S. Falke, M. Misera, U. Sterr, and C. Lisdat. Delivering pulsed and phase stable light to atoms of an optical clock. *Applied Physics B*, 107(2):301–311, 2012.

[81] Stephan Falke, Nathan Lemke, Christian Grebing, Burghard Lipphardt, Stefan Weyers, Vladislav Gerginov, Nils Huntemann, Christian Hagemann, Ali Al-Masoudi, Sebastian Häfner, Stefan Vogt, Uwe Sterr, and Christian Lisdat. A strontium lattice clock with 3×10^{17} inaccuracy and its frequency. *New Journal of Physics*, 16(7):073023, 2014.

[82] V. V. Flambaum, S. G. Porsev, and M. S. Safronova. Energy shift due to anisotropic blackbody radiation. *Phys. Rev. A*, 93:022508, Feb 2016.

[83] T. M. Fortier, N. Ashby, J. C. Bergquist, M. J. Delaney, S. A. Diddams, T. P. Heavner, L. Hollberg, W. M. Itano, S. R. Jefferts, K. Kim, F. Levi, L. Lorini, W. H. Oskay, T. E. Parker, J. Shirley, and J. E. Stalnaker. Precision atomic spectroscopy for improved limits on variation of the fine structure constant and local position invariance. *Phys. Rev. Lett.*, 98:070801, Feb 2007.

[84] R.W. Fox, C.W. Oates, and L.W. Hollberg. *Cavity-Enhanced Spectroscopies*, chapter Stabilizing diode lasers to high finesse cavities. Elsevier Science, New York, 2002.

[85] J Friebe, M Riedmann, T W?bbena, A Pape, H Kelkar, W Ertmer, O Terra, U Sterr, S Weyers, G Grosche, H Schnatz, and E M Rasel. Remote frequency measurement of the ${}^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_1$ transition in laser-cooled 24 Mg. *New Journal of Physics*, 13(12):125010, 2011.

[86] Kurt Gibble. Decoherence and collisional frequency shifts of trapped bosons and fermions. *Phys. Rev. Lett.*, 103:113202, Sep 2009.

[87] Kurt Gibble. Scattering of cold-atom coherences by hot atoms: Frequency shifts from background-gas collisions. *Phys. Rev. Lett.*, 110:180802, May 2013.

[88] L. I. R. Gil, R. Mukherjee, E. M. Bridge, M. P. A. Jones, and T. Pohl. Spin squeezing in a rydberg lattice clock. *Phys. Rev. Lett.*, 112:103601, Mar 2014.

[89] Patrick Gill. When should we change the definition of the second? *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 369(1953):4109–4130, 2011.

[90] E. G obel, I.M. Mills, and A.J. Wallard. *Le Système International d'unités*. Bureau International des poids et mesures, 8th edition, 2006.

[91] R. M. Godun, P. B. R. Nisbet-Jones, J. M. Jones, S. A. King, L. A. M. Johnson, H. S. Margolis, K. Szymaniec, S. N. Lea, K. Bongs, and P. Gill. Frequency Ratio of Two Optical Clock Transitions in $^{171}\mathrm{Yb}^{+}$ and Constraints on the Time Variation of Fundamental Constants. *Phys. Rev. Lett.*, 113:210801, Nov 2014.

[92] Christian Grebing, Ali Al-Masoudi, Sören Dörscher, Sebastian Häfner, Vladislav Gerginov, Stefan Weyers, Burghard Lipphardt, Fritz Riehle, Uwe Sterr, and Christian Lisdat. Realization of a timescale with an accurate optical lattice clock. *Optica*, 3(6):563–569, Jun 2016.

[93] O. W. Greenberg. *CPT* Violation Implies Violation of Lorentz Invariance. *Phys. Rev. Lett.*, 89:231602, Nov 2002.

[94] R. Grimm, M. Weidemüller, and Y.B. Ovchinnikov. Optical Dipole Traps for Neutral Atoms. In Benjamin Bederson and Herbert Walther, editors, *Advances in Atomic, Molecular, and Optical Physics*, volume 42 of *Advances In Atomic, Molecular, and Optical Physics*, pages 95 – 170. Academic Press, 2000.

[95] J. Guéna, M. Abgrall, D. Rovera, P. Laurent, B. Chupin, M. Lours, G. Santarelli, P. Rosenbusch, M. Tobar, R. Li, K. Gibble, A. Clairon, and S. Bize. Progress in Atomic Fountains at LNE-SYRTE. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 59:391–420, 2012.

[96] J. Guéna, M. Abgrall, D. Rovera, P. Rosenbusch, M. E. Tobar, Ph. Laurent, A. Clairon, and S. Bize. Improved tests of local position invariance using $^{87}\mathrm{Rb}$ and $^{133}\mathrm{Cs}$ fountains. *Phys. Rev. Lett.*, 109:080801, Aug 2012.

[97] J. R. Guest, N. D. Scielzo, I. Ahmad, K. Bailey, J. P. Greene, R. J. Holt, Z.-T. Lu, T. P. O'Connor, and D. H. Potterveld. Laser trapping of $^{225}\mathrm{Ra}$ and $^{226}\mathrm{Ra}$ with repumping by room-temperature blackbody radiation. *Phys. Rev. Lett.*, 98:093001, Feb 2007.

[98] H. Hachisu, K. Miyagishi, S.G. Porsev, A. Derevianko, V.D. Ovsiannikov, V.G. Pal'chikov, M. Takamoto, and H. Katori. Trapping of Neutral Mercury Atoms and Prospects for Optical Lattice Clocks. *Physical Review Letters*, 100(5):053001, February 2008.

[99] Sebastian Häfner, Stephan Falke, Christian Grebing, Stefan Vogt, Thomas Legero, Mikko Merimaa, Christian Lisdat, and Uwe Sterr. 8×10^{-17} fractional laser frequency instability with a long room-temperature cavity. *Opt. Lett.*, 40(9):2112–2115, May 2015.

[100] C. Hagemann, C. Grebing, T. Kessler, S. Falke, N. Lemke, C. Lisdat, H. Schnatz, F. Riehle, and U. Sterr. Providing 10^{-16} short-term stability of a 1.5- μm laser to optical clocks. *Instrumentation and Measurement, IEEE Transactions on*, 62(6):1556–1562, June 2013.

[101] Christian Hagemann, Christian Grebing, Christian Lisdat, Stephan Falke, Thomas Legero, Uwe Sterr, Fritz Riehle, Michael J. Martin, and Jun Ye. Ultrastable laser with average fractional frequency drift rate below 5×10^{-19}/s. *Opt. Lett.*, 39(17):5102–5105, Sep 2014.

[102] J.L. Hall, Wei Zhang, and Jun Ye. Accurate removal of RAM from FM laser beams. In *Frequency Control Symposium the European Frequency and Time Forum (FCS), 2015 Joint Conference of the IEEE International*, pages 713–716, April 2015.

[103] T.M. Hard. Laser Wavelength Selection and Output Coupling by a Grating. *Applied Optics*, 9(8):1825–1830, Aug 1970.

[104] Thomas P Heavner, Elizabeth A Donley, Filippo Levi, Giovanni Costanzo, Thomas E Parker, Jon H Shirley, Neil Ashby, Stephan Barlow, and S R Jefferts. First accuracy evaluation of NIST-F2. *Metrologia*, 51(3):174, 2014.

[105] C.H. Henry, N.A. Olsson, and N.K. Dutta. Locking range and stability of injection locked 1.54 μm ingaasp semiconductor lasers. *Quantum Electronics, IEEE Journal of*, 21(8):1152–1156, Aug 1985.

[106] I. R. Hill. *Development of an Apparatus for a Strontium Optical Lattice Optical Frequency Standard.* PhD thesis, Imperial College London, June 2012.

[107] Ian R Hill, Yuri B Ovchinnikov, Elizabeth M Bridge, E Anne Curtis, and Patrick Gill. Zeeman slowers for strontium based on permanent magnets. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 47(7):075006, 2014.

[108] I.R. Hill, Y.B. Ovchinnikov, E.M. Bridge, E.A. Curtis, S. Donnellan, and P. Gill. A simple, configurable, permanent magnet Zeeman slower for Sr. In *European Frequency and Time Forum (EFTF), 2012*, pages 545–549, April 2012.

[109] N. Hinkley, J. A. Sherman, N. B. Phillips, M. Schioppo, N. D. Lemke, K. Beloy, M. Pizzocaro, C. W. Oates, and A. D. Ludlow. An Atomic Clock with 10^{18} Instability. *Science*, 341(6151):1215–1218, 2013.

[110] R. Hobson, W. Bowden, S. A. King, P. E. G. Baird, I. R. Hill, and P. Gill. Modified hyper-ramsey methods for the elimination of probe shifts in optical clocks. *Phys. Rev. A*, 93:010501, Jan 2016.

[111] M. A. Hohensee, N. Leefer, D. Budker, C. Harabati, V. A. Dzuba, and V. V. Flambaum. Limits on violations of lorentz symmetry and the einstein equivalence principle using radio-frequency spectroscopy of atomic dysprosium. *Phys. Rev. Lett.*, 111:050401, Jul 2013.

[112] L Hollberg, S Diddams, A Bartels, T Fortier, and K Kim. The measurement of optical frequencies. *Metrologia*, 42(3):S105, 2005.

[113] T. Hong, C. Cramer, W. Nagourney, and E.N. Fortson. Optical Clocks Based on Ultranarrow Three-Photon Resonances in Alkaline Earth Atoms. *Physical Review Letters*, 94(5), February 2005.

[114] N. Huntemann. *High-Accuracy Optical Clock Based on the Octupole Transition in $^{171}Yb^{+}$*. PhD thesis, Universit at Hannover, 2014.

[115] N. Huntemann, B. Lipphardt, M. Okhapkin, Chr. Tamm, E. Peik, A. V. Taichenachev, and V. I. Yudin. Generalized ramsey excitation scheme with suppressed light shift. *Phys. Rev. Lett.*, 109:213002, Nov 2012.

[116] N. Huntemann, B. Lipphardt, Chr. Tamm, V. Gerginov, S. Weyers, and E. Peik. Improved Limit on a Temporal Variation of m_p/m_e from Comparisons of Yb^+ and Cs Atomic Clocks. *Phys. Rev. Lett.*, 113:210802, Nov 2014.

[117] N. Huntemann, M. Okhapkin, B. Lipphardt, S. Weyers, Chr. Tamm, and E. Peik. High-accuracy optical clock based on the octupole transition in $^{171}\mathrm{Yb}^+$. *Phys. Rev. Lett.*, 108:090801, Feb 2012.

[118] N. Huntemann, C. Sanner, B. Lipphardt, Chr. Tamm, and E. Peik. Single-ion atomic clock with 3×10^{-18} systematic uncertainty. *Phys. Rev. Lett.*, 116:063001, Feb 2016.

[119] Tetsuya Ido, Yoshitomo Isoya, and Hidetoshi Katori. Optical-dipole trapping of Sr atoms at a high phase-space density. *Phys. Rev. A*, 61:061403, May 2000.

[120] Hajime Inaba, Kazumoto Hosaka, Masami Yasuda, Yoshiaki Nakajima, Kana Iwakuni, Daisuke Akamatsu, Sho Okubo, Takuya Kohno, Atsushi Onae, and Feng-Lei Hong. Spectroscopy of ^{171}Yb in an optical lattice based on laser linewidth transfer using a narrow linewidth frequency comb. *Opt. Express*, 21(7):7891–7896, Apr 2013.

[121] W.M. Itano, J.C. Bergquist, J.J. Bollinger, J.M. Gilligan, D.J. Heinzen, F.L. Moore, M.G. Raizen, and D.J. Wineland. Quantum projection noise: Population fluctuations in two-level systems. *Physical Review A*, 47(5):3554–3570, May 1993.

[122] Herbert E. Ives and G. R. Stilwell. An experimental study of the rate of a moving atomic clock. *J. Opt. Soc. Am.*, 28(7):215–226, Jul 1938.

[123] D. Jennings, K. Evenson, and D.J.E. Knight. Optical frequency measurements. *Proceedings of the IEEE*, 74(1):168–179, Jan 1986.

[124] Y.Y. Jiang, A.D. Ludlow, N.D. Lemke, R.W. Fox, J.A. Sherman, L.-S. Ma, and C.W. Oates. Making optical atomic clocks more stable with 10^{-16}-level laser stabilization. *Nature Photonics*, 5(3):158–161, March 2011.

[125] L A M Johnson, P Gill, and H S Margolis. Evaluating the performance of the NPL femtosecond frequency combs: agreement at the 10 ^{21} level. *Metrologia*, 52(1):62, 2015.

[126] David J. Jones, Scott A. Diddams, Jinendra K. Ranka, Andrew Stentz, Robert S. Windeler, John L. Hall, and Steven T. Cundiff. Carrier-envelope phase control of femtosecond mode-locked lasers and direct optical frequency synthesis. *Science*, 288(5466):635–639, 2000.

[127] Nissim Kanekar. Constraining changes in the proton-electron mass ratio with inversion and rotational lines. *The Astrophysical Journal Letters*, 728(1):L12, 2011.

[128] H. Katori. Spectroscopy of strontium atoms in the Lamb-Dicke confinement. *Proceedings of the 6th Symposium on Frequency Standards and Metrology*, pages 323–330, 2002.

[129] H. Katori, T. Ido, Y. Isoya, and M. Kuwata-Gonokami. Magneto-Optical Trapping and Cooling of Strontium Atoms down to the Photon Recoil Temperature. *Physical Review Letters*, 82:1116–1119, Feb 1999.

[130] Hidetoshi Katori, Koji Hashiguchi, E. Yu. Il'inova, and V. D. Ovsiannikov. Magic wavelength to make optical lattice clocks insensitive to atomic motion. *Phys. Rev. Lett.*, 103:153004, Oct 2009.

[131] Hidetoshi Katori, V. D. Ovsiannikov, S. I. Marmo, and V. G. Palchikov. Strategies for reducing the light shift in atomic clocks. *Phys. Rev. A*, 91:052503, May 2015.

[132] Hidetoshi Katori, Tetsushi Takano, and Masao Takamoto. Optical lattice clocks and frequency comparison. *Journal of Physics: Conference Series*, 264(1):012011, 2011.

[133] Kaye and Laby. *Tables of Physical & Chemical Constants*, chapter 2.5.7 Refractive index of gases. NPL, 2015.

[134] Roy J. Kennedy and Edward M. Thorndike. Experimental establishment of the relativity of time. *Phys. Rev.*, 42:400–418, Nov 1932.

[135] T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M. J. Martin, L. Chen, and J. Ye. A sub-40-mHz-linewidth based on a silicon single-crystal optical cavity. *Nature Photonics*, 2012.

[136] Thomas Kessler, Thomas Legero, and Uwe Sterr. Thermal noise in optical cavities revisited. *J. Opt. Soc. Am. B*, 29(1):178–184, Jan 2012.

[137] S.A. King, R.M. Godun, S.A. Webster, H.S. Margolis, L.A.M. Johnson, K. Szymaniec, P.E.G. Baird, and P. Gill. Absolute frequency measurement of the $^2S_{1/2}$ – $^2F_{7/2}$ electric octupole transition in a single ion of $^{171}Yb^+$ with 10^{15} fractional uncertainty. *New Journal of Physics*, 14(1):013045, 2012.

[138] Masahiro Kitagawa and Masahito Ueda. Squeezed spin states. *Phys. Rev. A*, 47:5138–5143, Jun 1993.

[139] V. Alan Kostelecký and Stuart Samuel. Spontaneous breaking of Lorentz symmetry in string theory. *Phys. Rev. D*, 39:683–685, Jan 1989.

[140] G. Kramer, B. Lipphardt, and C.O. Weiss. Coherent frequency synthesis in the infrared. In *Frequency Control Symposium, 1992. 46th., Proceedings of the 1992 IEEE*, pages 39–43, May 1992.

[141] A. P. Kulosa, D. Fim, K. H. Zipfel, S. Rühmann, S. Sauer, N. Jha, K. Gibble, W. Ertmer, E. M. Rasel, M. S. Safronova, U. I. Safronova, and S. G. Porsev. Towards a Mg Lattice Clock: Observation of the $^1S_0$-$^3P_0$ Transition and Determination of the Magic Wavelength. *Phys. Rev. Lett.*, 115:240801, Dec 2015.

[142] A. Kuzmich, N. P. Bigelow, and L. Mandel. Atomic quantum non-demolition measurements and squeezing. *EPL (Europhysics Letters)*, 42(5):481, 1998.

[143] R. Le Targat. *Horloge à réseau optique au Strontium: une 2ème génération d'horloges à atatom froids*. PhD thesis, Laboratoire des systèmes de référence temps-espace, 2007.

[144] R. Le Targat, X. Baillard, M. Fouche, A. Brusch, O. Tcherbakoff, G.D. Rovera, and P. Lemonde. Accurate optical lattice clock with ^{87}Sr atoms. *Physical Review Letters*, 97(13):130801, September 2006.

[145] R. Le Targat, L. Lorini, Y. Le Coq, M. Zawada, J. Guéna, M. Abgrall, M. Gurov, P. Rosenbusch, D. G. Rovera, B. Nagórny, R. Gartman, P. G. Westergaard, M. E. Tobar, M. Lours, G. Santarelli, A. Clairon, S. Bize, P. Laurent, P. Lemonde, and J. Lodewyck. Experimental realization of an optical second with strontium lattice clocks. *Nat Commun*, 4, 07 2013.

[146] N. Leefer, C. T. M. Weber, A. Cingöz, J. R. Torgerson, and D. Budker. New limits on variation of the fine-structure constant using atomic dysprosium. *Phys. Rev. Lett.*, 111:060801, Aug 2013.

[147] T. Legero, T. Kessler, and U. Sterr. Tuning the thermal expansion properties of optical reference cavities with fused silica mirrors. *Journal of the Optical Society of America B*, 27(5):914–919, May 2010.

[148] T Legero, J S R Winfred, F Riehle, and U Sterr. Ultracold ^{88}Sr atoms for an optical lattice clock. In *Frequency Control Symposium, Joint with the 21st European Frequency and Time Forum. IEEE International*, pages 119–122, 2007.

[149] D.R. Leibrandt, J.C. Bergquist, and T. Rosenband. A cavity-stabilized laser with acceleration sensitivity below 10^{-12}/g. *ArXiv e-prints*, December 2012.

[150] D.R. Leibrandt, M.J. Thorpe, M. Notcutt, R.E. Drullinger, T. Rosenband, and J.C. Bergquist. Spherical reference cavities for frequency stabilization of lasers in non-laboratory environments. *Optics Express*, 19(4):3471–3482, Feb 2011.

[151] N. D. Lemke, J. von Stecher, J. A. Sherman, A. M. Rey, C. W. Oates, and A. D. Ludlow. p-wave cold collisions in an optical lattice clock. *Phys. Rev. Lett.*, 107:103902, Aug 2011.

[152] N.D. Lemke, A.D. Ludlow, Z.W. Barber, T.M. Fortier, S.A. Diddams, Y. Jiang, S.R. Jefferts, T.P. Heavner, T.E. Parker, and C.W. Oates. Spin-1/2 Optical Lattice Clock. *Physical Review Letters*, 103(6):063001, Aug 2009.

[153] P. Lemonde and P. Wolf. Optical lattice clock with atoms confined in a shallow trap. *Physical Review A*, 72(3):033409, Sep 2005.

[154] Y. Levin. Internal thermal noise in the ligo test masses: A direct approach. *Phys. Rev. D*, 57(2):659–662, 1998.

[155] Levshakov, S. A., Molaro, P., Lapinov, A. V., Reimers, D., Henkel, C., and Sakai, T. Searching for chameleon-like scalar fields with the ammonia method. *Astronomy and Astrophysics*, 512:A44, 2010.

[156] Liufeng Li, Fang Liu, Chun Wang, and Lisheng Chen. Measurement and control of residual amplitude modulation in optical phase modulation. *Review of Scientific Instruments*, 83(4):–, 2012.

[157] Ruoxin Li, Kurt Gibble, and Krzysztof Szymaniec. Improved accuracy of the NPL-CsF2 primary frequency standard: evaluation of distributed cavity phase and microwave lensing frequency shifts. *Metrologia*, 48(5):283, 2011.

[158] K.G. Libbrecht and J.L. Hall. A low-noise high-speed diode laser current controller. *Review of Scientific Instruments*, 64(8):2133–2135, 1993.

[159] Yi-Ge Lin, Qiang Wang, Ye Li, Fei Meng, Bai-Ke Lin, Er-Jun Zang, Zhen Sun, Fang Fang, Tian-Chu Li, and Zhan-Jun Fang. First Evaluation and Frequency Measurement of the Strontium Optical Lattice Clock at NIM. *Chinese Physics Letters*, 32(09):90601, 2015.

[160] C. Lisdat, G. Grosche, N. Quintin, C. Shi, S.M.F. Raupach, C. Grebing, D. Nicolodi, F. Stefani, A. Al-Masoudi, S. Drscher, S. Hfner, J.-L. Robyr, N. Chiodo, S. Bilicki, E. Bookjans, A. Koczwara, S. Koke, A. Kuhl, F. Wiotte, F. Meynadier, E. Camisard, M. Abgrall, M. Lours, T. Legero, H. Schnatz, U. Sterr, H. Denker, C. Chardonnet, Y. Le Coq, G. Santarelli, A. Amy-Klein, R. Le Targat, J. Lodewyck, O. Lopez, and P.-E. Pottie. A clock network for geodesy and fundamental science. *ArXiv e-prints*, November 2015.

[161] Ch. Lisdat, J. S. R. Vellore Winfred, T. Middelmann, F. Riehle, and U. Sterr. Collisional losses, decoherence, and frequency shifts in optical lattice clocks with bosons. *Phys. Rev. Lett.*, 103:090801, Aug 2009.

[162] G. Lochead, D. Boddy, D. P. Sadler, C. S. Adams, and M. P. A. Jones. Number-resolved imaging of excited-state atoms using a scanning autoionization microscope. *Phys. Rev. A*, 87:053409, May 2013.

[163] T.H. Loftus, T. Ido, M.M. Boyd, A.D. Ludlow, and Jun Ye. Narrow line cooling and momentum-space crystals. *Physical Review A*, 70:063413, Dec 2004.

[164] A.D. Ludlow. *The Strontium Optical Lattice Clock: Optical Spectroscopy with Sub-Hertz Accuracy.* PhD thesis, University of Colorado, 2008.

[165] A.D. Ludlow, M.M. Boyd, T. Zelevinsky, S.M. Foreman, S. Blatt, M. Notcutt, T. Ido, and Jun Ye. Systematic Study of the ^{87}Sr Clock Transition in an Optical Lattice. *Physical Review Letters*, 96(3), January 2006.

[166] A.D. Ludlow, X. Huang, M. Notcutt, T. Zanon-Willette, S.M. Foreman, M.M. Boyd, S. Blatt, and Jun Ye. Compact, thermal-noise-limited optical cavity for diode laser stabilization at 1×10^{-15}. *Optics Letters*, 32(6):641–3, March 2007.

[167] A.D. Ludlow, T. Zelevinsky, G.K. Campbell, S. Blatt, M.M. Boyd, M.H.G. de Miranda, M.J. Martin, J.W. Thomsen, S.M. Foreman, Jun Ye, T.M. Fortier, J.E. Stalnaker, S.A. Diddams, Y.Le Coq, Z.W. Barber, N. Poli, N.D. Lemke, K.M. Beck, and C.W. Oates. Sr lattice clock at 1x10^{-16} fractional uncertainty by remote optical evaluation with a Ca clock. *Science*, 319:1805, 2008.

[168] A.N. Luiten. Ongoing research on reference cavities at the University of Adelaide. For details see `http://www.adelaide.edu.au/ipas/research/nls/pmg-research/f-standards/`.

[169] A. Lurio, M. Mandel, and R. Novick. Second-Order Hyperfine and Zeeman Corrections for an (sl) Configuration. *Phys. Rev.*, 126:1758–1767, Jun 1962.

[170] L.-S. Ma, P. Jungner, Jun Ye, and J.L. Hall. Delivering the same optical frequency at two places: accurate cancellation of phase noise introduced by an optical fiber or other time-varying path. *Optics Letters*, 19(21):1777–9, November 1994.

[171] Long-Sheng Ma, Zhiyi Bi, Albrecht Bartels, Lennart Robertsson, Massimo Zucco, Robert S. Windeler, Guido Wilpers, Chris Oates, Leo Hollberg, and Scott A. Diddams. Optical frequency synthesis and comparison with uncertainty at the 10-19 level. *Science*, 303(5665):1843–1845, 2004.

[172] A.A. Madej, P. Dubé, Z. Zhou, and M. Gertsvolf. ^{88}sr^{+} 445-thz single-ion reference at the 10(-17) level via control and cancellation of systematic uncertainties and its measurement against the si second. *Phys. Rev. Lett.*, 109(20), November 2012.

[173] H.S. Margolis, G.P. Barwood, G. Huang, H.A. Klein, S.N. Lea, K. Szymaniec, and P. Gill. Hertz-level measurement of the optical clock frequency in a single ^{88}Sr^{+} ion. *Science*, 306(5700):1355–8, November 2004.

[174] B. J. Martin. *Building a Better Atomic Clock*. PhD thesis, University of Colorado, 2014.

[175] M. J. Martin. *Quantum Metrology and Many-Body Physics: Pushing the Frontier of the Optical Lattice Clock*. PhD thesis, University of Colorado, 2013.

[176] M. J. Martin, M. Bishof, M. D. Swallows, X. Zhang, C. Benko, J. von Stecher, A. V. Gorshkov, A. M. Rey, and Jun Ye. A quantum many-body spin system in an optical lattice clock. *Science*, 341(6146):632–636, 2013.

[177] M. J. Martin, D. Meiser, J.W. Thomsen, J. Ye, and M. J. Holland. Extreme nonlinear response of ultranarrow optical transitions in cavity QED for laser stabilisation. *Phys. Rev. A*, 84(063813), 2011.

[178] J. J. McFerran, L. Yi, S. Mejri, S. Di Manno, W. Zhang, J. Guéna, Y. Le Coq, and S. Bize. Neutral atom frequency reference in the deep ultraviolet with fractional uncertainty = 5.7×10^{-15}. *Phys. Rev. Lett.*, 108:183004, May 2012.

[179] T. E. Mehlstäubler, K. Moldenhauer, M. Riedmann, N. Rehbein, J. Friebe, E. M. Rasel, and W. Ertmer. Observation of sub-doppler temperatures in bosonic magnesium. *Phys. Rev. A*, 77:021402, Feb 2008.

[180] D. Meiser, J. Ye, D. R. Carlson, and M. J. Holland. Prospects for a milihertz-linewidth laser. *Phys. Rev. Lett.*, 102, 2009.

[181] D. Meiser, J. Ye, and M. J. Holland. Spin squeezing in optical lattice clocks via lattice-based qnd measurements. *New Journal of Physics*, 10(073014), 2008.

[182] A. A. Michelson and E. W. Morley. On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, Series 3 Vol. 34(203):333–345, 1887.

[183] P.G. Mickelson, Y.N. Martinez de Escobar, P. Anzel, B.J. DeSalvo, S.B. Nagel, A.J. Traverso, M. Yan, and T.C. Killian. Repumping and spectroscopy of laser-cooled Sr atoms using the (5s5p) ^{3}P$_2$ - (5s4d) ^{3}D$_2$ transition. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 42(23):235001, 2009.

[184] T. Middelmann, C. Lisdat, S. Falke, J.S.R.V. Winfred, F. Riehle, and U. Sterr. Tackling the Blackbody Shift in a Strontium Optical Lattice Clock. *Instrumentation and Measurement, IEEE Transactions on*, 60(7):2550 –2557, July 2011.

[185] Thomas Middelmann, Stephan Falke, Christian Lisdat, and Uwe Sterr. High accuracy correction of blackbody radiation shift in an optical lattice clock. *Phys. Rev. Lett.*, 109:263004, Dec 2012.

[186] J. Millen, G. Lochead, and M. P. A. Jones. Two-Electron Excitation of an Interacting Cold Rydberg Gas. *Phys. Rev. Lett.*, 105:213004, Nov 2010.

[187] J. Mitroy. Polarizabilities of the beryllium clock transition. *Phys. Rev. A*, 82:052516, Nov 2010.

[188] J. Mitroy and J.Y. Zhang. Dispersion and polarization interactions of the strontium atom. *Molecular Physics*, 108(15):1999–2006, 2010.

[189] Takashi Mukaiyama, Hidetoshi Katori, Tetsuya Ido, Ying Li, and Makoto Kuwata-Gonokami. Recoil-limited laser cooling of ^{87}Sr atoms near the fermi temperature. *Phys. Rev. Lett.*, 90:113002, Mar 2003.

[190] S. Nagel, C. Simien, S. Laha, P. Gupta, V. Ashoka, and T. Killian. Magnetic trapping of metastable $^3\mathrm{P}_2$ atomic strontium. *Physical Review A*, 67(1):1–4, January 2003.

[191] Yoshiaki Nakajima, Hajime Inaba, Kazumoto Hosaka, Kaoru Minoshima, Atsushi Onae, Masami Yasuda, Takuya Kohno, Sakae Kawato, Takao Kobayashi, Toshio Katsuyama, and Feng-Lei Hong. A multi-branch, fiber-based frequency comb with millihertz-level relative linewidths using an intra-cavity electro-optic modulator. *Opt. Express*, 18(2):1667–1676, Jan 2010.

[192] T. L. Nicholson. *A new record in atomic clock performance*. PhD thesis, University of Colorado, 2015.

[193] T. L. Nicholson, S. L. Campbell, R. B. Hutson, G. E. Marti, B. J. Bloom, R. L. McNally, W. Zhang, M. D. Barrett, M. S. Safronova, G. F. Strouse, W. L. Tew, and J. Ye. Systematic evaluation of an atomic clock at 2×10^{-18} total uncertainty. *Nat Commun*, 6, 04 2015.

[194] T.L. Nicholson, M. J. Martin, J. R. Williams, B. J. Bloom, M. Bishof, M.D. Swallows, S. L. Campbell, and J. Ye. Comparison of Two Independent Sr Optical Clocks with 1×10^{-17} Stability at 10^3 s. *Physical Review Letters*, 109(230801), December 2012.

[195] Daniele Nicolodi, Berengere Argence, Wei Zhang, Rodolphe Le Targat, Giorgio Santarelli, and Yann Le Coq. Spectral purity transfer between optical wavelengths at the 10^{-18} level. *Nat Photon*, 8(3):219–223, Mar 2014. Letter.

[196] M. Niering, R. Holzwarth, J. Reichert, P. Pokasov, Th. Udem, M. Weitz, T. W. Hänsch, P. Lemonde, G. Santarelli, M. Abgrall, P. Laurent, C. Salomon, and A. Clairon. Measurement of the hydrogen $1S$- $2S$ transition frequency by phase coherent comparison with a microwave cesium fountain clock. *Phys. Rev. Lett.*, 84:5496–5499, Jun 2000.

[197] L. Nilse, H.J. Davies, and C.S. Adams. Synchronous Tuning of Extended Cavity Diode Lasers: The Case for an Optimum Pivot Point. *Applied Optics*, 38(3):548–553, Jan 1999.

[198] M. A. Norcia and J. K. Thompson. A Cold-Strontium Laser in the Superradiant Crossover Regime. *ArXiv e-prints*, October 2015.

[199] M. A. Norcia, M. N. Winchester, J. R. K. Cline, and J. Thompson. Superradiance on the millihertz linewidth strontium clock transition. *ArXiv e-prints*, 2016.

[200] Dennis Normile and Daniel Clery. First global telescope opens an eye on the cold universe. *Science*, 333(6051):1820–1823, 2011.

[201] M. Notcutt, L.-S. Ma, Jun Ye, and J.L. Hall. Simple and compact 1-Hz laser system via an improved mounting configuration of a reference cavity. *Optics Letters*, 30(14):1815–7, July 2005.

[202] K. Numata, A. Kemery, and J. Camp. Thermal-Noise Limit in the Frequency Stabilization of Lasers with Rigid Cavities. *Physical Review Letters*, 93(25):1–4, December 2004.

[203] Kenji Numata, Giuseppe Bertolotto Bianc, Naoko Ohishi, Atsushi Sekiya, Shigemi Otsuka, Keita Kawabe, Masaki Ando, and Kimio Tsubono. Measurement of the intrinsic mechanical loss of low-loss samples using a nodal support. *Physics Letters A*, 276(14):37 – 46, 2000.

[204] S. Origlia et al. Development of a strontium optical lattice clock for the SOC mission on the ISS. *ArXiv e-prints*, March 2016.

[205] W. Oskay, S. Diddams, E. Donley, T. Fortier, T. Heavner, L. Hollberg, W. Itano, S. Jefferts, M. Delaney, K. Kim, F. Levi, T. Parker, and J. Bergquist. Single-Atom Optical Clock with High Accuracy. *Physical Review Letters*, 97(2), July 2006.

[206] Y.B. Ovchinnikov. A Zeeman slower based on magnetic dipoles. *Optics Communications*, 276(2):261–267, August 2007.

[207] V. D. Ovsiannikov, V. G. Pal'chikov, A. V. Taichenachev, V. I. Yudin, and Hidetoshi Katori. Multipole, nonlinear, and anharmonic uncertainties of clocks of Sr atoms in an optical lattice. *Phys. Rev. A*, 88:013405, Jul 2013.

[208] Vitaly Ovsiannikov, Vitaly Pal'chikov, Alexey Taichenachev, Valeriy Yudin, Hidetoshi Katori, and Masao Takamoto. Magic-wave-induced $^1S_0$ - $^3P_0$ transition in even isotopes of alkaline-earth-metal-like atoms. *Physical Review A*, 75(2), February 2007.

[209] Ekkehard Peik, Tobias Schneider, and Christian Tamm. Laser frequency stabilization to a single ion. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 39(1):145, 2006.

[210] Steven Peil, Scott Crane, James L. Hanssen, Thomas B. Swanson, and Christopher R. Ekstrom. Tests of local position invariance using continuously running atomic clocks. *Phys. Rev. A*, 87:010102, Jan 2013.

[211] N. Poli, Z.W. Barber, N.D. Lemke, C.W. Oates, L.S. Ma, J.E. Stalnaker, T.M. Fortier, S.A. Diddams, L. Hollberg, J.C. Bergquist, A. Brusch, S. Jefferts, T. Heavner, and T. Parker. Frequency evaluation of the doubly forbidden $^1S_0$ – $^3P_0$ transition in bosonic ^{174}Yb. *Physical Review A*, 77(5):050501, 2008.

[212] N. Poli, M. Schioppo, S. Vogt, St. Falke, U. Sterr, Ch. Lisdat, and G.M. Tino. A transportable strontium optical lattice clock. *Applied Physics B*, 117(4):1107–1116, 2014.

[213] S.G. Porsev and A. Derevianko. Hyperfine quenching of the metastable $^3P_{0,2}$ states in divalent atoms. *Physical Review A*, 69(4):042506, Apr 2004.

[214] S.G. Porsev and A. Derevianko. Multipolar theory of blackbody radiation shift of atomic energy levels and its implications for optical lattice clocks. *Physical Review A*, 74(2):4–7, August 2006.

[215] S.G. Porsev, A.D. Ludlow, M.M. Boyd, and Jun Ye. Determination of Sr properties for a high-accuracy optical clock. *Physical Review A*, 78(3):032508, Sep 2008.

[216] T. Pruttivarasin, M. Ramm, S. G. Porsev, I. I. Tupitsyn, M. S. Safronova, M. A. Hohensee, and H. Haffner. Michelson-morley analogue for electrons using trapped ions to test lorentz symmetry. *Nature*, 517(7536):592–595, Jan 2015. Letter.

[217] E.L. Raab, M. Prentiss, A. Cable, S. Chu, and D.E. Pritchard. Trapping of Neutral Sodium Atoms with Radiation Pressure. *Physical Review Letters*, 59:2631–2634, Dec 1987.

[218] N.F. Ramsey. A Molecular Beam Resonance Method with Separated Oscillating Fields. *Physical Review*, 78(6):695–699, June 1950.

[219] J. Reichert, M. Niering, R. Holzwarth, M. Weitz, Th. Udem, and T. W. Hänsch. Phase coherent vacuum-ultraviolet to radio frequency comparison with a mode-locked laser. *Phys. Rev. Lett.*, 84:3232–3235, Apr 2000.

[220] Sascha Reinhardt, Guido Saathoff, Henrik Buhr, Lars A. Carlson, Andreas Wolf, Dirk Schwalm, Sergei Karpuk, Christian Novotny, Gerhard Huber, Marcus Zimmermann, Ronald Holzwarth, Thomas Udem, Theodor W. Hansch, and Gerald Gwinner. Test of relativistic time dilation with fast optical atomic clocks at different velocities. *Nat Phys*, 3(12):861–864, Dec 2007.

[221] A.M. Rey, A.V. Gorshkov, C.V. Kraus, M.J. Martin, M. Bishof, M.D. Swallows, X. Zhang, C. Benko, J. Ye, N.D. Lemke, and A.D. Ludlow. Probing many-body interactions in an optical lattice clock. *Annals of Physics*, 340(1):311 – 351, 2014.

[222] H. P. Robertson. Postulate *versus* observation in the special theory of relativity. *Rev. Mod. Phys.*, 21:378–382, Jul 1949.

[223] T. Rosenband, D.B. Hume, P.O. Schmidt, C.W. Chou, A. Brusch, L. Lorini, W.H. Oskay, R.E. Drullinger, T.M. Fortier, J.E. Stalnaker, S.A. Diddams, W.C. Swann, N.R. Newbury, W.M. Itano, D.J. Wineland, and J.C. Bergquist. Frequency Ratio of Al^+ and Hg^+ Single-Ion Optical Clocks; Metrology at the 17th Decimal Place. *Science*, 319(5871):1808–1812, March 2008.

[224] T. Rosenband, P.O. Schmidt, D.B. Hume, W.M. Itano, T.M. Fortier, J.E. Stalnaker, K. Kim, S.A. Diddams, J.C.J. Koelemeij, J.C. Bergquist, and D.J. Wineland. Observation of the $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_0$ Clock Transition in $^{27}\mathrm{Al}^+$. *Physical Review Letters*, 98(22):220801, May 2007.

[225] M. A. Rowe, D. Kielpinski, V. Meyer, C. A. Sackett, W. M. Itano, C. Monroe, and D. J. Wineland. Experimental violation of a bell's inequality with efficient detection. *Nature*, 409(6822):791–794, Feb 2001.

[226] J. Rutman and F.L. Walls. Characterization of frequency stability in precision frequency sources. *Proceedings of the IEEE*, pages 952–960, July 1991.

[227] M. S. Safronova, V. A. Dzuba, V. V. Flambaum, U. I. Safronova, S. G. Porsev, and M. G. Kozlov. Atomic properties of Cd-like and Sn-like ions for the development of frequency standards and search for the variation of the fine-structure constant. *Phys. Rev. A*, 90:052509, Nov 2014.

[228] M. S. Safronova, V. A. Dzuba, V. V. Flambaum, U. I. Safronova, S. G. Porsev, and M. G. Kozlov. Highly charged Ag-like and In-like ions for the development of atomic clocks and the search for α variation. *Phys. Rev. A*, 90:042513, Oct 2014.

[229] M. S. Safronova, V. A. Dzuba, V. V. Flambaum, U. I. Safronova, S. G. Porsev, and M. G. Kozlov. Highly charged ions for atomic clocks, quantum information, and search for α variation. *Phys. Rev. Lett.*, 113:030801, Jul 2014.

[230] M. S. Safronova, S. G. Porsev, U. I. Safronova, M. G. Kozlov, and Charles W. Clark. Blackbody-radiation shift in the Sr optical atomic clock. *Phys. Rev. A*, 87:012509, Jan 2013.

[231] B.E.A. Saleh and M.C. Teich. *Fundamentals of Photonics*. Wiley Series in Pure and Applied Optics. Wiley, 2007.

[232] R. Santra, K. Christ, and C.H. Greene. Properties of metastable alkaline-earth-metal atoms calculated using an accurate effective core potential. *Phys. Rev. A*, 69(4), 2004.

[233] Robin Santra, Ennio Arimondo, Tetsuya Ido, Chris Greene, and Jun Ye. High-Accuracy Optical Clock via Three-Level Coherence in Neutral Bosonic ^{88}Sr. *Physical Review Letters*, 94(17), May 2005.

[234] Juna Sathian and Esa Jaatinen. Intensity dependent residual amplitude modulation in electro-optic phase modulators. *Appl. Opt.*, 51(16):3684–3691, Jun 2012.

[235] S. Schiller, A. Gorlitz, A. Nevsky, S. Alighanbari, S. Vasilyev, C. Abou-Jaoudeh, G. Mura, T. Franzen, U. Sterr, S. Falke, C. Lisdat, E. Rasel, A. Kulosa, S. Bize, J. Lodewyck, G.M. Tino, N. Poli, M. Schioppo, K. Bongs, Y. Singh, P. Gill, G. Barwood, Y. Ovchinnikov, J. Stuhler, W. Kaenders, C. Braxmaier, R. Holzwarth, A. Donati, S. Lecomte, D. Calonico, and F. Levi. The space optical clocks project: Development of high-performance transportable and breadboard optical clocks and advanced subsystems. In *European Frequency and Time Forum (EFTF), 2012*, pages 412–418, April 2012.

[236] S. Schiller, G.M. Tino, P. Gill, C. Salomon, U. Sterr, E. Peik, A. Nevsky, A. Grlitz, D. Svehla, G. Ferrari, N. Poli, L. Lusanna, H. Klein, H. Margolis, P. Lemonde, P. Laurent, G. Santarelli, A. Clairon, W. Ertmer, E. Rasel,

J. Mller, L. Iorio, C. Lmmerzahl, H. Dittus, E. Gill, M. Rothacher, F. Flechner, U. Schreiber, V. Flambaum, Wei-Tou Ni, Liang Liu, Xuzong Chen, Jingbiao Chen, Kelin Gao, L. Cacciapuoti, R. Holzwarth, M.P. He, and W. Schfer. Einstein gravity explorera medium-class fundamental physics mission. *Experimental Astronomy*, 23(2):573–610, 2009.

[237] L. Schmger, O. O. Versolato, M. Schwarz, M. Kohnen, A. Windberger, B. Piest, S. Feuchtenbeiner, J. Pedregosa-Gutierrez, T. Leopold, P. Micke, A. K. Hansen, T. M. Baumann, M. Drewsen, J. Ullrich, P. O. Schmidt, and J. R. Crespo Lpez-Urrutia. Coulomb crystallization of highly charged ions. *Science*, 347(6227):1233–1236, 2015.

[238] H. Schnatz, B. Lipphardt, J. Helmcke, F. Riehle, and G. Zinner. First phase-coherent frequency measurement of visible radiation. *Phys. Rev. Lett.*, 76:18–21, Jan 1996.

[239] J. A. Sherman, N. D. Lemke, N. Hinkley, M. Pizzocaro, R. W. Fox, A. D. Ludlow, and C. W. Oates. High-accuracy measurement of atomic polarizability in an optical lattice clock. *Phys. Rev. Lett.*, 108:153002, Apr 2012.

[240] N Shiga and M Takeuchi. Locking the local oscillator phase to the atomic phase via weak measurement. *New Journal of Physics*, 14(2):023034, 2012.

[241] P. T. Starkey, C. J. Billington, S. P. Johnstone, M. Jasperse, K. Helmerson, L. D. Turner, and R. P. Anderson. A scripted control system for autonomous hardware-timed experiments. *Review of Scientific Instruments*, 84(8), 2013.

[242] Daniel A. Steck. *Quantum and Atom Optics*. Online, http://steck.us/teaching, 0.8.3 edition, 2012.

[243] S. Stellmer. *Degenerate quantum gases of strontium*. PhD thesis, University of Innsbruck, 2013.

[244] Simon Stellmer, Rudolf Grimm, and Florian Schreck. Production of quantum-degenerate strontium gases. *Phys. Rev. A*, 87:013611, Jan 2013.

[245] Simon Stellmer, Benjamin Pasquiou, Rudolf Grimm, and Florian Schreck. Laser cooling to quantum degeneracy. *Phys. Rev. Lett.*, 110:263003, Jun 2013.

[246] Simon Stellmer and Florian Schreck. Reservoir spectroscopy of $5s5p\ ^3P_2$ - $5snd\ ^3D_{1,2,3}$ transitions in strontium. *Phys. Rev. A*, 90:022512, Aug 2014.

[247] Simon Stellmer, Meng Khoon Tey, Bo Huang, Rudolf Grimm, and Florian Schreck. Bose-Einstein Condensation of Strontium. *Physical Review Letters*, 103(20):200401, Nov 2009.

[248] Jörn Stenger, Christian Tamm, Nils Haverkamp, Stefan Weyers, and Harald R. Telle. Absolute frequency measurement of the 435.5 nm 171Yb$^+$-clock transition with a Kerr-lens mode-locked femtosecond laser. *Opt. Lett.*, 26(20):1589–1591, Oct 2001.

[249] U. Sterr, C. Degenhardt, H. Stoehr, Ch. Lisdat, H. Schnatz, J. Helmcke, F. Riehle, G. Wilpers, Ch. Oates, and L. Hollberg. The optical calcium frequency standards of {PTB} and {NIST}. *Comptes Rendus Physique*, 5(8):845 – 855, 2004. Fundamental metrology.

[250] S.A. Strelkin, A.A. Galyshev, O.I. Berdasov, A.Yu. Gribov, D.V. Sutyrin, K.Yu. Khabarova, N.N. Kolachevsky, and S.N. Slyusarev. Narrow line cooling of 88sr atoms in the magneto-optical trap for precision frequency standard. *Physics Procedia*, 72:184 – 188, 2015. Conference of Physics of Nonequilibrium Atomic Systems and Composites, {PNASC} 2015, 18-20 February 2015 and Conference of Heterostructures for Microwave, Power and Optoelectronics: Physics, Technology and Devices, 19 February 2015.

[251] M. D. Swallows, M. J. Martin, M. Bishof, C. Benko, Y. Lin, S. Blatt, A. M. Rey, and J. Ye. Operating a ^{87}sr optical lattice clock with high precision and at high density. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, 59(3), March 2012.

[252] Matthew D. Swallows, Michael Bishof, Yige Lin, Sebastian Blatt, Michael J. Martin, Ana Maria Rey, and Jun Ye. Suppression of collisional shifts in a strongly interacting lattice clock. *Science*, 331(6020):1043–1046, 2011.

[253] A. V. Taichenachev, V. I. Yudin, and C. W. Oates. Optical lattice polarization effects on magnetically induced optical atomic clock transitions. *Phys. Rev. A*, 76:023806, Aug 2007.

[254] A. V. Taichenachev, V. I. Yudin, C. W. Oates, C. W. Hoyt, Z. W. Barber, and L. Hollberg. Magnetic field-induced spectroscopy of forbidden optical transitions with application to lattice-based optical atomic clocks. *Phys. Rev. Lett.*, 96:083001, Mar 2006.

[255] A.V. Taichenachev, V.I. Yudin, C.W. Oates, Z.W. Barber, N.D. Lemke, A.D. Ludlow, U. Sterr, Ch. Lisdat, and F. Riehle. Compensation of field-induced frequency shifts in Ramsey spectroscopy of optical clock transitions. *JETP Letters*, 90(11):713–717, 2010.

[256] A.V. Taichenachev, V.I. Yudin, V.D. Ovsiannikov, and V.G. Pal'chikov. Optical Lattice Polarization Effects on Hyperpolarizability of Atomic Clock Transitions. *Physical Review Letters*, 97(17):173601, 2006.

[257] A.V. Taichenachev, V.I. Yudin, V.D. Ovsiannikov, V.G. Pal'chikov, and C.W. Oates. Frequency Shifts in an Optical Lattice Clock Due to Magnetic-Dipole and Electric-Quadrupole Transitions. *Physical Review Letters*, 101(19):193601, Nov 2008.

[258] M. Takamoto, F.-L. Hong, R. Higashi, Y. Fujii, M. Imae, and H. Katori. Improved Frequency Measurement of a One-Dimensional Optical Lattice Clock with a Spin-Polarized Fermionic ^{87}Sr Isotope. *Journal of the Physical Society of Japan*, 75(10):104302, 2006.

[259] M. Takamoto, F.-L. Hong, R. Higashi, and H. Katori. An optical lattice clock. *Nature*, 435(7040):321–324, May 2005.

[260] M. Takamoto, H. Katori, S.I. Marmo, V.D. Ovsiannikov, and V.G. Pal'chikov. Prospects for Optical Clocks with a Blue-Detuned Lattice. *Physical Review Letters*, 102(6):63002, 2009.

[261] Masao Takamoto, Tetsushi Takano, and Hidetoshi Katori. Frequency comparison of optical lattice clocks beyond the Dick limit. *Nat Photon*, 5(5):288–292, May 2011.

[262] C.T. Taylor, M. Notcutt, E.K. Wong, A.G. Mann, and D.G. Blair. Measurement of the coefficient of thermal expansion of a cryogenic, all-sapphire, Fabry-Perot optical cavity. *Optics Communications*, 131(46):311 – 314, 1996.

[263] H. R. Telle, D. Meschede, and T. W. Hänsch. Realization of a new concept for visible frequency division: phase locking of harmonic and sum frequencies. *Opt. Lett.*, 15(10):532–534, May 1990.

[264] H.R. Telle, B. Lipphardt, and J. Stenger. Kerr-lens, mode-locked lasers as transfer oscillators for optical frequency measurements. *Applied Physics B*, 74(1):1–6, 2002.

[265] M.J. Thorpe, D.R. Leibrandt, T.M. Fortier, and T. Rosenband. Measurement and real-time cancellation of vibration-induced phase noise in a cavity-stabilized laser. *Optics Express*, 18(18):18744–18751, Aug 2010.

[266] M.J. Thorpe, L. Rippe, T. M. Fortier, M. S. Kirchner, and T. Rosenband. Frequency stabilization to 6×10^{-16} via spectral-hole burning. *Nature Photonics*, 5(11):688–693, 2011.

[267] S. Truppe, R. J. Hendricks, S. K. Tokunaga, H. J. Lewandowski, M. G. Kozlov, Christian Henkel, E. A. Hinds, and M. R. Tarbutt. A search for varying fundamental constants using hertz-level frequency measurements of cold CH molecules. *Nat Commun*, 4, Oct 2013. Article.

[268] R. Tyumenev, M. Favier, S. Bilicki, E. Bookjans, R. Le Targat, J. Lodewyck, D. Nicolodi, Y. Le Coq, M. Abgrall, Guèna, L. De Sarlo, and S. Bize. Comparing a mercury optical lattice clock with microwave and optical frequency standards. *ArXiv e-prints*, Mar 2016.

[269] Th. Udem, S. A. Diddams, K. R. Vogel, C. W. Oates, E. A. Curtis, W. D. Lee, W. M. Itano, R. E. Drullinger, J. C. Bergquist, and L. Hollberg. Absolute Frequency Measurements of the Hg^{+} and Ca Optical Clock Transitions with a Femtosecond Laser. *Phys. Rev. Lett.*, 86:4996–4999, May 2001.

[270] Th. Udem, J. Reichert, R. Holzwarth, and T. W. Hänsch. Absolute Optical Frequency Measurement of the Cesium D_1 Line with a Mode-Locked Laser. *Phys. Rev. Lett.*, 82:3568–3571, May 1999.

[271] Ichiro Ushijima, Masao Takamoto, Manoj Das, Takuya Ohkubo, and Hidetoshi Katori. Cryogenic optical lattice clocks. *Nat Photon*, 9(3):185–189, 03 2015.

[272] Jean-Philippe Uzan. The fundamental constants and their variation: observational and theoretical status. *Rev. Mod. Phys.*, 75:403–455, Apr 2003.

[273] Ken Van Tilburg, Nathan Leefer, Lykourgos Bougas, and Dmitry Budker. Search for ultralight scalar dark matter with atomic spectroscopy. *Phys. Rev. Lett.*, 115:011802, Jun 2015.

[274] H. Wallis and W. Ertmer. Broadband laser cooling on narrow transitions. *J. Opt. Soc. Am. B*, 6(11):2211–2219, Nov 1989.

[275] G. Wang and A. Ye. Possibility of using Zn as the quantum absorber for a laser-cooled neutral atomic optical frequency standard. *Physical Review A*, 76(4):043409, Oct 2007.

[276] J. K. Webb, J. A. King, M. T. Murphy, V. V. Flambaum, R. F. Carswell, and M. B. Bainbridge. Indications of a spatial variation of the fine structure constant. *Phys. Rev. Lett.*, 107:191101, Oct 2011.

[277] J. K. Webb, M. T. Murphy, V. V. Flambaum, V. A. Dzuba, J. D. Barrow, C. W. Churchill, J. X. Prochaska, and A. M. Wolfe. Further evidence for cosmological evolution of the fine structure constant. *Phys. Rev. Lett.*, 87:091301, Aug 2001.

[278] S.A. Webster and P. Gill. Force-insensitive optical cavity. *Optics Letters*, 2011.

[279] S.A. Webster, M. Oxborrow, S. Pugla, J. Millo, and P. Gill. Thermal-noise-limited optical cavity. *Physical Review A*, 77(3), March 2008.

[280] John Weiner, Vanderlei S. Bagnato, Sergio Zilio, and Paul S. Julienne. Experiments and theory in cold and ultracold collisions. *Rev. Mod. Phys.*, 71:1–85, Jan 1999.

[281] Jonathan D. Weinstein, Kyle Beloy, and Andrei Derevianko. Entangling the lattice clock: Towards heisenberg-limited timekeeping. *Phys. Rev. A*, 81:030302, Mar 2010.

[282] P.G. Westergaard. *Horloge à r eseau optique au Strontium : en quete de la performance ultime*. PhD thesis, l'Edite de Paris, 2010.

[283] P.G. Westergaard, J. Lodewyck, and P. Lemonde. Minimizing the Dick effect in an optical lattice clock. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 57(3):623 –628, march 2010.

[284] P.G. Westergaard, J. Lodewyck, L. Lorini, A. Lecallier, E.A. Burt, M. Zawada, J. Millo, and P. Lemonde. Lattice-Induced Frequency Shifts in Sr Optical Lattice Clocks at the 10^{-17} Level. *Physical Review Letters*, 106(21):210801, May 2011.

[285] Philip G. Westergaard, R. Christensen, Bjarke T. David Tieri, Rastin Matin, John Cooper, Murray Holland, Jun Ye, and Jan W. Thomsen. Observation of motion-dependent nonlinear dispersion with narrow-linewidth atoms in an optical cavity. *Phys. Rev. Lett.*, 114:093002, Mar 2015.

[286] Guy K. White. Reference materials for thermal expansion: certified or not? *Thermochimica Acta*, 218:83 – 99, 1993.

[287] D. J. Wineland, J. J. Bollinger, W. M. Itano, F. L. Moore, and D. J. Heinzen. Spin squeezing and reduced quantum noise in spectroscopy. *Phys. Rev. A*, 46:R6797–R6800, Dec 1992.

[288] D.J. Wineland and W.M. Itano. Laser cooling of atoms. *Phys. Rev. A*, 20(4):1521–1540, October 1979.

[289] P. Wolf, Ch.J. Bord, A. Clairon, L. Duchayne, A. Landragin, P. Lemonde, G. Santarelli, W. Ertmer, E. Rasel, F.S. Cataliotti, M. Inguscio, G.M. Tino, P. Gill, H. Klein, S. Reynaud, C. Salomon, E. Peik, O. Bertolami, P. Gil, J. Pramos, C. Jentsch, U. Johann, A. Rathke, P. Bouyer, L. Cacciapuoti, D. Izzo, P. De Natale, B. Christophe, P. Touboul, S.G. Turyshev, J. Anderson, M.E. Tobar, F. Schmidt-Kaler, J. Vigu, A.A. Madej, L. Marmet, M.-C. Angonin, P. Delva, P. Tourrenc, G. Metris, H. Mller, R. Walsworth, Z.H. Lu, L.J. Wang, K. Bongs, A. Toncelli, M. Tonelli, H. Dittus, C. Lmmerzahl, G. Galzerano, P. Laporta, J. Laskar, A. Fienga, F. Roques, and K. Sengstock. Quantum physics exploring gravity in the outer solar system: the SAGAS project. *Experimental Astronomy*, 23(2):651–687, 2009.

[290] P. Wolf and G. Petit. Relativistic theory for clock syntonization and the realization of geocentric coordinate times. *Astronomy & Astrophsics*, 304:653, December 1995.

[291] N.C. Wong and J.L. Hall. Servo control of amplitude modulation in frequency-modulation spectroscopy: demonstration of shot-noise-limited detection. *Journal of the Optical Society of America B*, 2(9):1527–1533, Sep 1985.

[292] X. Xu, T.H. Loftus, J.L. Hall, A. Gallagher, and Jun Ye. Cooling and trapping of atomic strontium. *Journal of the Optical Society of America B (Optical Physics)*, 20(5):968 – 76, 2003.

[293] Atsushi Yamaguchi, Miho Fujieda, Motohiro Kumagai, Hidekazu Hachisu, Shigeo Nagano, Ying Li, Tetsuya Ido, Tetsushi Takano, Masao Takamoto, and Hidetoshi Katori. Direct comparison of distant optical lattice clocks at the 10^{-16} uncertainty. *Applied Physics Express*, 4(8):082203, 2011.

[294] Kazuhiro Yamanaka, Noriaki Ohmae, Ichiro Ushijima, Masao Takamoto, and Hidetoshi Katori. Frequency Ratio of ^{199}Hg and ^{87}Sr Optical Lattice Clocks beyond the SI Limit. *Phys. Rev. Lett.*, 114:230801, Jun 2015.

[295] Anpei Ye and Guangfu Wang. Dipole polarizabilities of $ns^2\,{}^1S_0$ and $nsnp\,{}^3P_0$ states and relevant magic wavelengths of group-iib atoms. *Phys. Rev. A*, 78:014502, Jul 2008.

[296] V. I. Yudin, A. V. Taichenachev, and M. Yu. Basalaev. Synthetic Frequency Protocol in the Ramsey Spectroscopy of Clock Transitions. *ArXiv e-prints*, Jan 2016.

[297] V. I. Yudin, A. V. Taichenachev, C. W. Oates, Z. W. Barber, N. D. Lemke, A. D. Ludlow, U. Sterr, Ch. Lisdat, and F. Riehle. Hyper-Ramsey spectroscopy of optical clock transitions. *Phys. Rev. A*, 82:011804, Jul 2010.

[298] T. Zanon-Willette, S. Almonacil, E. de Clercq, A. D. Ludlow, and E. Arimondo. Quantum engineering of atomic phase shifts in optical clocks. *Phys. Rev. A*, 90:053427, Nov 2014.

[299] T. Zanon-Willette, E. de Clercq, and E. Arimondo. Probe light-shift elimination in Generalized Hyper-Ramsey quantum clocks. *ArXiv e-prints*, Nov 2015.

[300] T. Zanon-Willette, V. I. Yudin, and A. V. Taichenachev. Generalized hyper-Ramsey resonance with separated oscillating fields. *Phys. Rev. A*, 92:023416, Aug 2015.

[301] Thomas Zanon-Willette, Andrew D. Ludlow, Sebastian Blatt, Martin M. Boyd, Ennio Arimondo, and Jun Ye. Cancellation of Stark Shifts in Optical Lattice Clocks by Use of Pulsed Raman and Electromagnetically Induced Transparency Techniques. *Phys. Rev. Lett.*, 97:233001, Dec 2006.

[302] T. Zelevinsky, S. Kotochigova, and Jun Ye. Precision test of mass-ratio variations with lattice-confined ultracold molecules. *Phys. Rev. Lett.*, 100:043201, Jan 2008.

[303] W. Zhang, M. J. Martin, C. Benko, J. L. Hall, J. Ye, C. Hagemann, T. Legero, U. Sterr, F. Riehle, G. D. Cole, and M. Aspelmeyer. Reduction of residual amplitude modulation to 1×10^{-6} for frequency modulation and laser stabilization. *Opt. Lett.*, 39(7):1980–1983, Apr 2014.

[304] X. Zhang, M. Bishof, S. L. Bromley, C. V. Kraus, M. S. Safronova, P. Zoller, A. M. Rey, and J. Ye. Spectroscopic observation of SU(N)-symmetric interactions in Sr orbital magnetism. *Science*, 345(6203):1467–1473, 2014.