Jason Hogan Lecture Questions

Questions marked in green were answered during the Q&A session. Original questions listed without correction for grammar/spelling. Where a slide number was given it is shown.

Q1 [slide 3] Is it possible to use atom interferometry to cover the part of the spectrum between the PTA and space-based arrays as well?

Probably not, at least with the current approach. The low frequency cutoff of the sensitivity is set by the duration of the atom interferometer sequence. We typically operate the atom interferometer as an accelerometer, so if the gravitation wave does not change appreciably ("accelerate") during the interferometer, the response is reduced. This leads to a roll off proportional to frequency squared at low frequencies, below a fraction of a Hz. While it is possible to imagine lowering the frequency cutoff somewhat, a practical limit (due to atom loss from imperfect vacuum) is probably 10 to 100 seconds. So it is possible to access some of the mHz frequency range, but below that is unlikely to work.

Q2 [slide 5] When the two detections are, say, six months apart, what tells us that they came from the same source? And we are not "pointing" using two unrelated signals?

Yes, the idea is to observe the same signal at multiple points in the orbit, so if there is confusion between two sources then this could be confounding. However, it is possible to avoid this in many cases. First of all, assuming that both of the signals in question are gravitational waves, then this is actually a somewhat "good" problem to have, since it implies that the detector has access to a lot of potential science signals! Nevertheless, it is important to be able to resolve the individual sources for many reasons. Ultimately this is a question about data analysis, and there has been much work done in this area by the gravitational wave community (in particular for LISA) that we can leverage. In fact, the number of simultaneous sources is expected to be much higher in LISA, so there the problem of resolving individual sources is even more acute. I can't do justice to the analysis techniques here, but I will just point out that there are usually ways to tell simultaneous sources apart, assuming there are not too many. For example, the sources typically will have different orbital frequencies, chirp rates, etc., since they are from totally different astrophysical systems. So they can be separately resolved in the Fourier domain, for instance. I will also note that the strategy of detecting a long-lived source is (ideally) to continuously monitor the signal over, say, six months, rather than make two measurements six months apart. This way it is possible to monitor the evolution of the signal (or multiple simultaneous signals) as the Earth goes around Sun, and use the observed variations to constrain the location(s).

Q3 [slide 7] Where are the ports on the cartoon atom interferometer on slide 7? Thanks!

The ports are on the right. The atom begins on the left, and the end of the sequence (final beamsplitter) is on the right. The two ports are in different internal atom states (red vs. blue) and are also traveling at different speeds (slope of lines), so they can be resolved either with a spin-dependent measurement or by their spatial separation.

Q4 [slide 9] How do you control whether absorption or stimulated emission happens, when light is shined upon the system?

Assuming a coherent process (negligible spontaneous emission), the evolution of the atom state in the presence of the light is to smoothly oscillate between the ground and excited states as shown by the plot of state populations versus time. This is known as Rabi oscillations. The evolution of the wavefunction is deterministic under the atom-light interaction Hamiltonian. So we can choose the final state of the system simply by choosing the duration of the laser pulse, since this determines how many Rabi oscillations the system undergoes. In fact, we also control the strength of the atom-light interaction itself (via the intensity of the laser), and this determines the Rabi frequency (how fast the state oscillates between ground and excited). So experimentally what we need to control is the product of the Rabi frequency and the pulse duration, which we call the "pulse area".

Q5 [slide 8-10] How do we give exactly half the atoms a kicked/increased velocity?

It is best to think about how the atom interferometer works for a single atom. While we actually use an ensemble of perhaps 10^6 atoms, the interference should be understood on a single atom basis. So it's not right to say that a we give half the atoms a kick. Instead, we take a single atom and divide its wavefunction quantum mechanically. In this language, "half" of this single atom's wavefunction remains in the initial state (slow), while the other half gets a kick. More precisely, we put the single atom into a superposition of two states: one which absorbed a photon and one which did not. This means that the interferometer diagram on slide 10 should be understood as the paths that a single atom takes during the sequence. This atom simultaneously follows two paths (the upper path and the lower path), and when we complete the interferometer we have to add these two amplitudes together to determine the probability of finding the atom in one of the two output ports. The contributions from the two paths can interfere constructively or destructively, depending on the phase of the interferometer. Now when we consider the ensemble of 10^6 atoms, we simply have the same interference experiment happening for all the atoms in parallel, with each atom in a superposition of states and subsequently interfering. We like to say that each atom can only interfere with itself, which is generally true in quantum mechanics.

Q6 [slide 10] What are the output ports of the interferometer?

The output ports are the two paths on the right of the (diamond shaped) interferometer diagram. Time runs from left to right, and the final beamsplitter occurs at the end of the sequence on the far right. After the final beamsplitter, there are two paths (one red and one blue), and these are the output ports. The probability of finding an atom in one of these ports depends on the phase difference between the two paths in the interferometer. To measure this probability, we can wait some small amount of time for the two paths to spatially separate (they have different velocities since their slopes are different on the diagram). When we take a picture of the interference pattern at the end of the sequence, we will in general see two clumps of atoms, corresponding to the two output ports. We estimate the probability (and from this infer the phase) by counting the number of atoms in the two ports.

Q7 [slide 11] Approximately how long are the beamsplitter/mirror pulses? How does it compare to the total travel time?

Approximately 10 to 100 microseconds typically. The total travel time is usually several hundred milliseconds, or even up to several seconds for a 10-meter scale atomic fountain.

Q8 Can an atomic interferrormeter work in a continues mode, as the atoms act as Beamsplitters or Mirror based on its phase.

The experiments we do are usually pulsed, rather than continuous, since we start off by preparing a distinct ensemble of cold atoms and then use that for the interferometer. So there is some experimental cycle where we first prepare the atoms, then do the interferometry, and finally do the atom detection. The repetition time of the experiment is roughly on the order of 1 Hz (sometimes slower for more elaborate cooling protocols). However, there are alternative approaches to atom interferometry that can be realized in a continuous mode. In this case, instead of distinct ensembles, a continuous atomic beam is usually used. This allows for a continuous monitoring of the interferometer phase and is particularly useful in applications where high bandwidth are important, such as when implementing atom interferometer gyroscopes for inertial navigation. Disadvantages of beam sources include that it is difficult to cool the atoms as well as with distinct ensembles, and there are some constraints on the interferometer geometry.

Q9 How long baseline will be required for an atomic interferrometer to detect Gravitational Waves ? Will its sensitivity be a function of frequency as in case of LIGO

We expect that a kilometer-scale baseline would be required. We are investigating several possible sites with existing kilometer-long vertical shafts that could be used. Yes, the sensitivity is generally a function of frequency. At frequencies below the inverse duration of the atom interferometer, the sensitivity falls off, typically proportional to frequency squared. At high frequencies we expect to be limited by the repetition rate of the experiment (Nyquist rate), which is likely a few Hz.

Q10 [slide 19] what would NOT be a good inertial test mass?

A "good" inertial test mass could be considered any object that ideally follows its geodesic as determined by gravity, and that is not strongly perturbed by non-gravitational forces. This includes forces from electric and magnetic fields, as well as vibration forces transmitted to the object by whatever is supporting it against gravity (assuming it is on Earth). One reason why atoms are good inertial test masses is because their susceptibility to electromagnetic forces is suppressed, since they are electrically neutral. Furthermore, the remaining sensitivity to electromagnetic fields that the atom does possess is well understood. All atoms of a particular species are identical, so we can measure the electromagnetic couplings precisely, and we can be sure that these measurements will correspond to the atoms used as test masses. This can be contrasted with the situation for macroscopic inertial test masses that must be built by people. In this case, each test mass is a unique artifact and must be individually characterized. Finally, atoms are good test masses because it is practical to put them into free-fall during the inertial measurements. The atoms we use for interferometry are allowed to fall under gravity while we measure their acceleration, and this substantially suppresses the influence of local seismic noise, for example. By comparison, the inertial test masses of LIGO must be isolated from vibration noise using sophisticated vibration isolation systems. Otherwise, the vibration noise will perturb the test masses away from their nominal geodesics, and this will obscure the gravitational wave measurement. In the case of atom interferometry, the equivalent test masses are the atoms themselves, and we allow the atoms to fall freely during the measurement. This means the atom test masses are decoupled from the vibration noise even without a vibration isolation system. For obvious reasons, it is not practical to drop the LIGO end mirrors in free-fall, so a vibration isolation system is needed.

Q11 [slide 19] Is it the same laser going back and forth 7 times?

We have one laser on each end of the experiment, one sending light upwards and one sending light downwards. The 7 distinct pulses shown in this figure on slide 19 are sent from alternating direction from the lasers at two locations (one at x=0 and the other at x=L). The different colors of the photons (light gray and dark gray) indicate the direction of the pulse and therefore which laser it came from. In practice, we often use one physical laser for both of these sets of pulses, and we use a mirror to change the direction of some of the pulses to realize light propagating in either direction, as needed.

Q12 How would you do sky localization with this technology?

By observing a gravitational wave signal for an extended period of time, it is possible to infer sky position information by incorporating our knowledge of the motion of the Earth around the

Sun and accounting for how this motion should affect the received signal. For example, as the Earth moves in its orbit around the Sun, the total distance between the source and the detector may modulate slightly. This variation in distance will affect the arrival time of the signal at the detector, so the Earth's motion can cause phase delays in the measured gravitational signal. The more the Earth moves, the larger the phase modulation, so it is beneficial to observe the same gravitational wave source for a substantial fraction of a year to maximize the Earth's motion. It is also possible to infer information about the polarization of a gravitational wave using extended observation, even with a single detector with only a single baseline. For observing polarization, as the Earth spins the detector orientation changes with respect to the gravitational wave signal, and this causes a modulation in the signal strength (depending on the projection of the polarization onto the detector baseline). So extended observation allows us to learn about sky position and polarization even with just a single detector. We can think of this as using the multiple observations at different times to effectively synthesize a multiple detector array, in this case with detectors separated in time rather than in space.

Q13 [slide 21] How will MAGIS-100 mitigate the effect of gravitational gradient with depth ?

Gravitational gradient noise (GGN) is likely the dominate noise source limiting any terrestrial gravitation wave detector at low frequencies. MAGIS-100 is not expected to be limited by gravity gradient noise because, as it is a demonstrator, the instrument is not sensitive enough to detect gravitational waves. However, we anticipate that a full-scale version of MAGIS would be limited by GGN. In MAGIS-100, we are exploring possible approaches to mitigate GGN. For example, by using more than two measurement points, it is in principle possible to measure and subtract GGN in some cases. GGN is due to the seismic motion of the Earth near the experiment, and so generally there is position variation in the strength of the gravitation forces exerted by these local masses (since gravity varies like distance squared from the source). By comparison, a gravitational wave is basically a perfect a plane wave, so the effect of the gravitational wave strain leads to a purely linear response across the baseline. By using more than two atom interferometers spread out along the baseline, we can look for any potential non-linear variation with position that would arise from GGN from local masses. In MAGIS-100 we plan to test this idea by using three atom sources.

Q14 [slide 21] What kinds of sources would MAGIS-100 be expected to detect?

MAGIS-100 is not expected to be sensitive enough to detect gravitational waves. There are several ultralight (wave-like) dark matter candidates that MAGIS-100 can search for, including both scalar- and vector-coupled models. For a full-scale detector based on the MAGIS concept, we expect to be able to see many of the same kinds of sources that LIGO has detected, such as binary black holes and neutron stars. In addition, it may be possible to see some sources that do not make it to the LIGO frequency band because they merge at lower peak orbital frequencies (for example, white dwarf binaries or black hole binaries in a different mass range). There is also always the possibility of something unexpected whenever you look in an unexplored frequency range.

Q15 [slide 22] Why do you need three atom sources in MAGIS-100?

See Q13 response. We plan on using the three atom sources to explore cancelation strategies for gravity gradient noise by attempting to measure curvature of the response across the baseline.

Q16 On the slide showing time evolution of binary mergers in strain-frequency space, they each had a kink where they turn downward at higher frequency - what is the cause of this? why do they start at low strain?

The kink is the result of an assumption about how long the signals can be observed. The detectability tends to be higher at lower frequencies because the signals last for a long time, so it is possible to average the data to reduce noise. This explains why the strain plots go up at first as you go to lower frequencies for a given source. At lower frequencies it is easier to see the same source because you can average for more periods. However, this trend does not go on forever, because we do not collect data forever. So even if a particular low frequency source would spend 100 years emitting gravitational waves near some frequency (for example), we will not be able to average the signal for this long because of the finite lifetime of the instrument. Here we assume a total data acquisition time of one year. The kink occurs right at the point when the lifetime of the source (at that frequency) exceeds the one year maximum averaging time. Below that we can no longer benefit from the increased source lifetime.

Q17 What do we mean by Rabi frequency ?

The Rabi frequency is a measure of the strength of the coupling between light and the atom. It is proportional to the strength of the electric field of the light and the dipole moment of the atom for the transition in question. It is called the Rabi "frequency" because this coupling strength determines the rate that the atom will oscillate between the ground and excited states during the light pulse.

Q18 If gravitational wave signals can be measured over a longer time, could there come up problems with some kind of pileup between several sources?

Yes, this can be an issue, but it's a great problem to have because it means there is a lot of great gravitational wave science to be done! Pileup is more likely at lower frequencies where there are more sources, so it is a particular issue for LISA. As it's in the middle, the mid-band range is likely to have more pileup concerns than LIGO, but less than LISA. There has already been substantial effort in the gravitational wave community to develop analysis techniques to address pileup in LISA, and detectors in the mid-band could likely benefit from this work as well.