

Environmental Noise in the Virgo Interferometer

Casey Stark
EGO Summer Student 2008
University of Southern California

August 5, 2008

1 The Virgo Interferometer

The Virgo Interferometer is a 3 km Fabry-Perot interferometer based in Cascina, Italy at the European Gravitational Observatory Labs. The interferometer was constructed to measure the strain of the distance between reference masses in order to measure gravitational waves. For reference, the interferometer was built to destructively interfere, so that the output is always dark, which is why we call the output signal the dark fringe signal.

Most of the work at Virgo is aimed at increasing measurement sensitivity. Currently, the biggest obstacle to gravitational wave observation is the required precision of the instruments. There are so many kinds of noise that keep the interferometer from working at the optimal sensitivity. However, by investigating noise in the device, understanding it, and then taking steps to mitigate it, eventually we should be able to use the Virgo Interferometer for gravitational wave observation.

2 The Effects of the Environment

The near perfection of the instruments inside this interferometer is astounding, but alone, it is not enough to negate the effects of the environment on our readings. Considering only the limitations of physics, the Virgo Interferometer should be much, much more sensitive than it is now. Unfortunately, we cannot isolate the interferometer from the outside world. It will always be on Earth, contacting the ground somehow, inside a building with other electronic devices, and be surrounded by noisy machines.

But this doesn't mean that we can't always do a better job. Gravitational wave observation is still a very young field and is a long way from being useful. If you look closely at the sensitivity curves of any of the modern gravitational interferometers, you will see that it is riddled with environmental noise. Using the noise projections of physical limitations, we should have a very smooth, simple sensitivity curve (basically the design curve), but in the real curve, we see many, many peaks that shouldn't be there. This is environmental noise and it is everywhere in these interferometers.

2.1 Types of Environmental Noise

There are many ways to classify environmental noise, but the common types are seismic, acoustic, magnetic, and electric.

Seismic noise is generated by anything that can shake the ground. This means the Earth itself and anything with moving parts near the interferometer. Any engines, pumps, and fans in the interferometer

buildings, cars, or foot steps can be picked up by the interferometer. Seismic noise is the most common problem at lower frequencies, but it is one of the easiest to mitigate. Usually it only requires some extra padding or modifying the resonances of structures.

Acoustic noise on the other hand, is the hardest to get rid of. Anything that shakes the interferometer through sound waves (instead of through the ground) is called acoustic noise. Basically, this means everything that makes noise: fans, pumps, planes, thunder, etc.

Magnetic noise enters the interferometer by coupling with the magnets on the payloads in each tower. Almost everything with an EMF can generate magnetic noise, but usually only large devices can create fields strong enough to affect the interferometer.

3 The 41-45 Hz Forest

On April 30th, 2008, J. Marque first made note of noise in the 41-45 Hz range [1]. We decided to investigate this forest so we could find the sources and begin to mitigate this noise.

3.1 Survey of Central Building

The first task to perform was to look at the dark fringe signal itself and check it against environmental sensors to see if there is any correlation. This preliminary survey would help give us a better idea of how this noise is created and picked up by the interferometer. Unfortunately, the interferometer was off at this time.

On May 8th, there was a disaster at the North End building. Around 8am that day, one of the viewports on the tower cracked and the tower imploded. Luckily, the North End vacuum seal was closed, so only the North End tower was compromised. A lot of the North End equipment has to be replaced (the mirror, viewports, etc.) or at least cleaned due to contamination. Currently, people from every department in Virgo are working hard to recover from this setback and to prepare the interferometer for when it is scheduled to be back online in September.

This left us with a big problem. It is very difficult to investigate noise in a signal that is off. To work around this, we looked at dark fringe data before the incident to create a reference for this noise. Using data during the very last science mode run (when the interferometer is locked and the data is good for gravitational wave analysis), we looked at the spectrum of every episensor and microphone in the central building and their coherence with the dark fringe. See Figure 1.

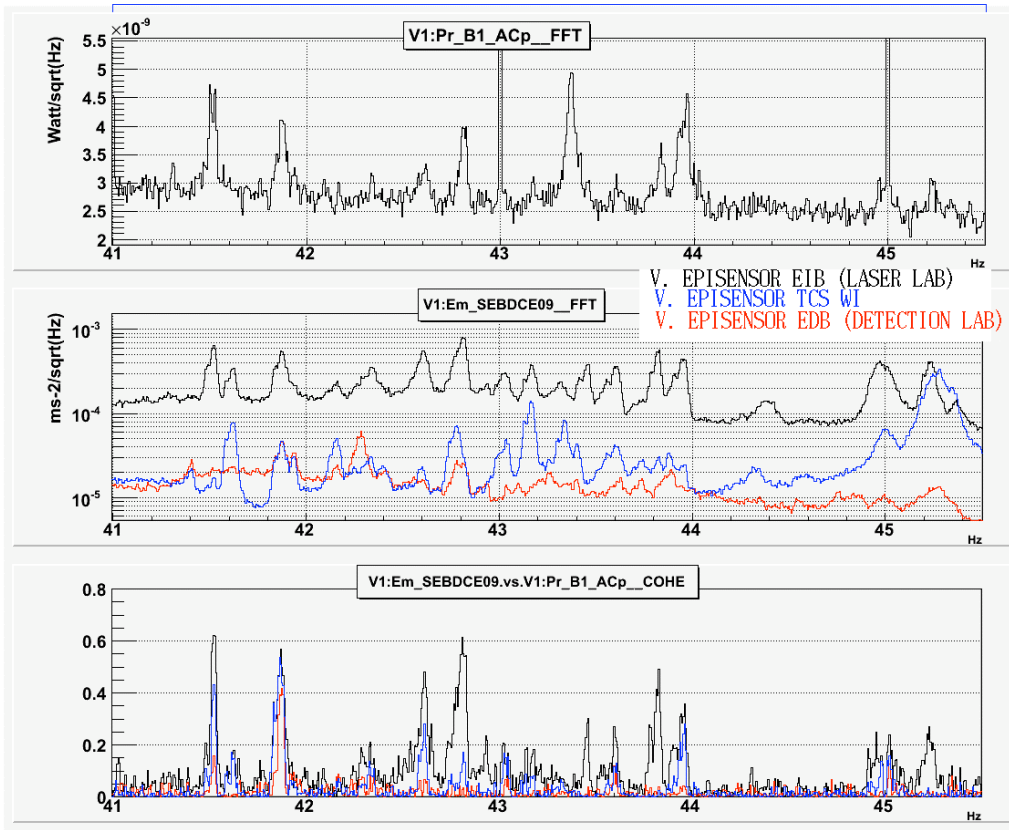


Figure 1: Spectra of the dark fringe, various episensors, and their coherence.

Using plots like these, we picked the most prominent peaks in the forest (those that showed up across multiple sensors and had the highest amplitude and coherence) to look for. These peaks are at the frequencies 41.5, 41.8, 42.6, 42.8, 43.4, 44.0, 45.0, and 45.2 Hz.

The most important thing that Figure 1 shows us is that this noise is best heard in the laser lab, where the external injection bench (EIB) is. We checked every episensor channel in the central building and the episensor mounted on the EIB was the loudest and most coherent with the dark fringe at every peak in the forest except 41.8 and 45.0 Hz. This evidence suggested that this noise enters the interferometer through the external injection bench.

We already had evidence that the EIB has a broad vertical resonance stretching from 30-50 Hz [2]. If there are sources of constant frequency noise in this range, they should excite the EIB enough for the noise to enter the interferometer and to be seen in the dark fringe. We can't be sure that this is the exact coupling path, but considering these two facts, it seems highly likely.

Another early assumption in our investigation was that the sources of all of these peaks are fans in the central building. This is based on the fact that the forest is just below the main line frequency (50 Hz) and fans should run just a bit slower than that because of losses to friction. This was also likely because we see almost all of the peaks in the forest in both seismic and acoustic sensors. We were sure that all of these fans are in the central building as well because the North and West end buildings and the more cleaner building episensors showed nothing in this range.

3.2 Probe Tests

The next step in our investigation was to use seismic probes (some portable accelerometers) to measure the fan noise directly. Because we were assuming that the noise is entering the interferometer through the EIB, we decided to measure the racks in the laser lab and platform racks and turbo pumps close to the laser lab. See Figure 2.




Figure 2: Map of the central building, labeled with the towers close to the laser lab. INJ is injection, PR is power recycling, BS is beam splitter, and WI is west input.

We mounted the accelerometers vertically on the racks, because the peaks were largest in vertical episensor channels, and horizontally on the turbo pumps, because it was the only way to mount the probes close to the cooling fans. We also left the probes overnight, so we could use the data from midnight (UTC), when the readings would not be disturbed by people, cars, etc.

The only problem with probing at this time is that we cannot compare the probe noise directly to the dark fringe because the interferometer is off. This made it very difficult to determine what noise would be picked up by the interferometer and to definitely identify the sources. To work around this, we compared the probe data against the fixed episensors (usually SEBDCE09, the vertical channel of the EIB episensor), using previous interferometer data as a reference for the forest. The episensor signal acted as a crude substitute for the dark fringe signal in this way.

While we were looking through the probe data, my mentor and I outlined a basic method for the noise hunt. To be a candidate for a certain peak in the forest, the source we were measuring with the probe had two meet 2 conditions.

First, there had to be strong coherence at that frequency between the probe and the fixed episensor. This tells us that the probe and the episensor are picking up the same signal at that frequency. Of course, this doesn't guarantee that the source the probe is measuring is the original source of this peak. It could easily pick up noise from elsewhere just like the fixed episensor.

Therefore, we also looked at the strength of the peak in the probe channel second. If it was much stronger than the peak in the episensor channel, then the source we were measuring is a good candidate for this peak. In cases where there were several candidates, we assumed that the source with the strongest peak was the original source.

Our procedure here wasn't very exact, but it is only because of our limitations. With the interferometer being off, having only two accelerometers (meaning two measurements per day), and the inherent difficulty in determining the noise coupling path, it is no wonder that we could only do so much with the probes. If the interferometer had been on and the peaks did not shift (we noticed a variance of about 0.1 Hz using the fixed episensors), this would have been a lot easier.

3.2.1 Preliminary Results

A couple of things were obvious after only a few days of probing. The first was that the turbo pump cooling fans definitely create the 45.0 and 45.2 Hz peaks. All of the turbo pump probes picked up two very strong peaks (because there are two cooling fans) around 45 Hz and were coherent with the central building episensors. See Figure 3.

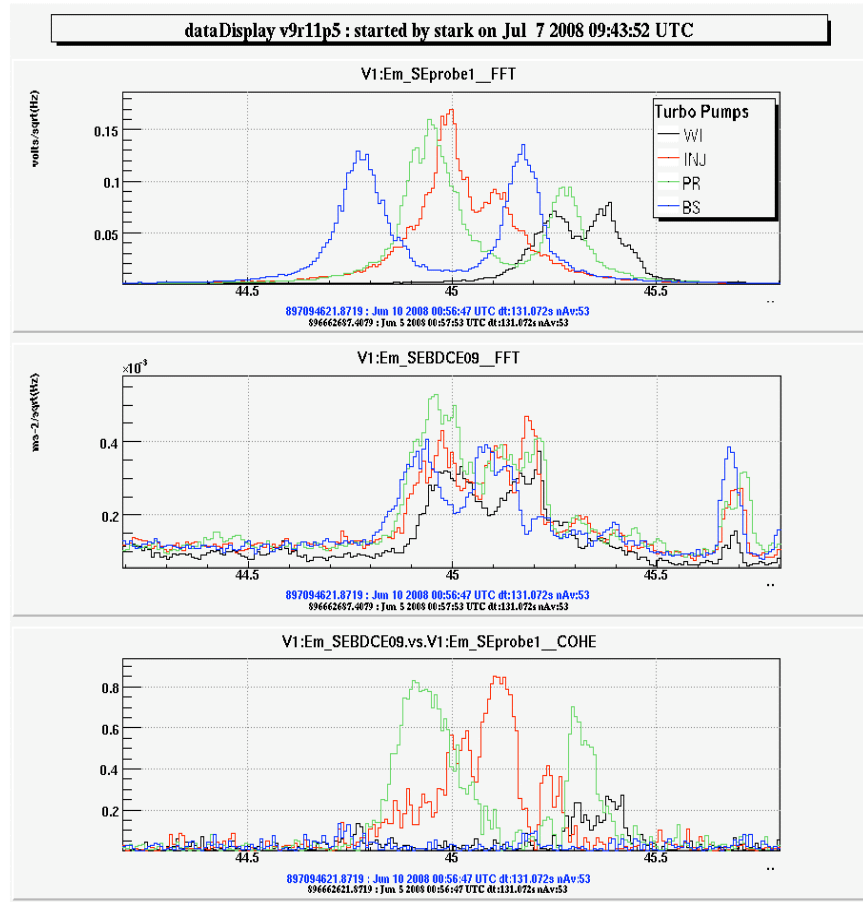


Figure 3: Spectra of four turbo pump cooling fans.

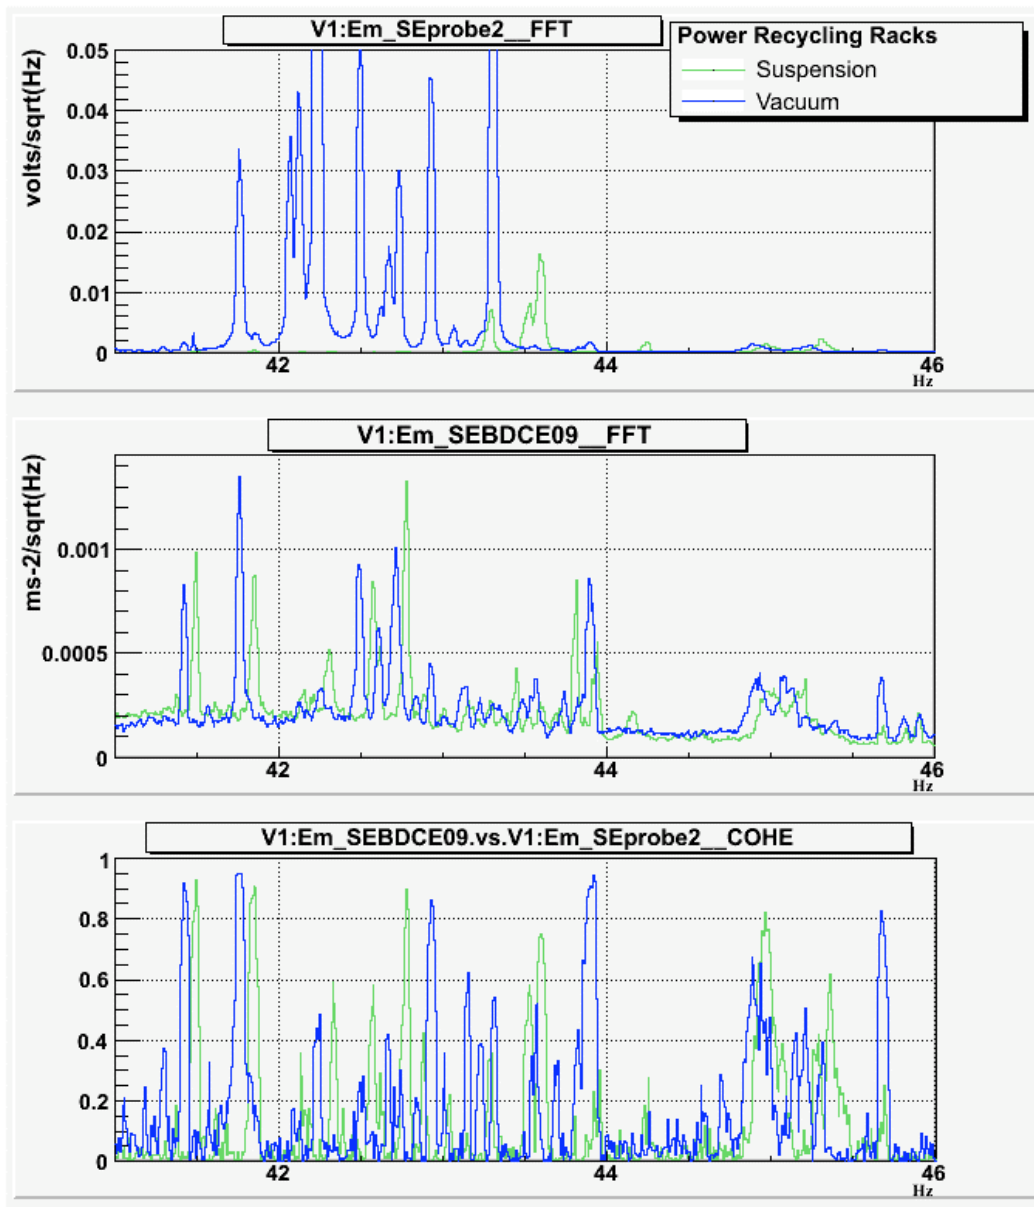


Figure 4: Comparison of vacuum and suspension racks. Similar coherence, but the suspension rack signal is so weak that it can't be generating any peaks in the EIB episensor.

Another preliminary result was that the vacuum racks are much noisier than the suspension racks. See Figure 4. Although the coherence patterns are very similar, the probe shows that the suspension rack is much quieter, meaning that the strong coherence is only present because the suspension rack feels the noise of the vacuum rack.

One surprising result we discovered was that the racks in the laser lab don't create much noise. In the beginning, we guessed that they would create most of this forest, but the probe data suggested otherwise. For some reason, the laser lab episensor doesn't pick up most of the strongest peaks generated by those racks even though they are only a few meters away. See Figure 5. The one exception to this is the mysterious 41.8 Hz peak, which we will discuss later.

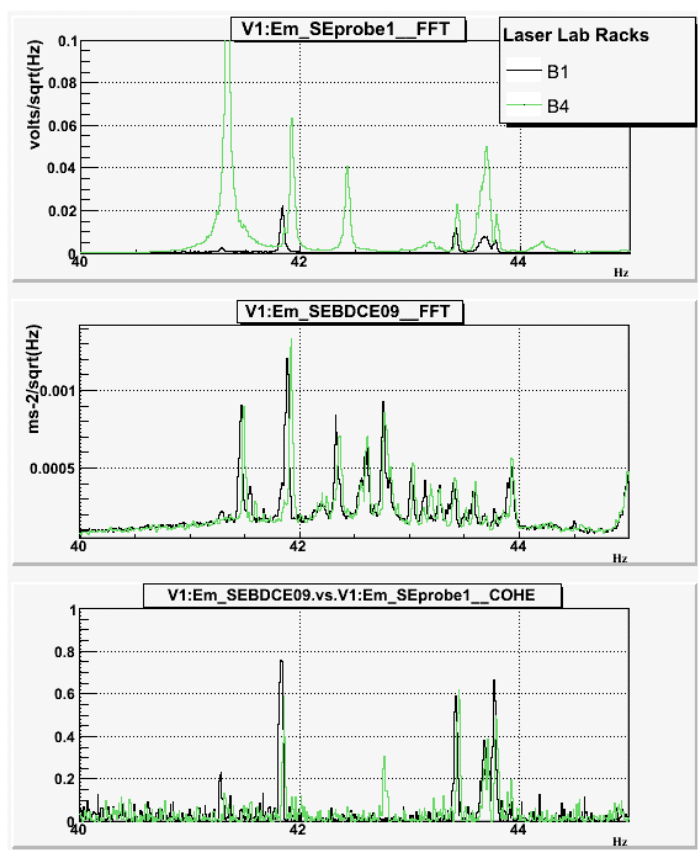


Figure 5: Spectra of laser lab rack noise.

Figure 6.

We had hoped for better results, but it was still pleasing to see several of the peaks disappear. Peaks at 41.5, 41.8, 42.3, and 42.8 Hz disappeared and almost all of the 45 Hz hump vanished. Unfortunately, we did not shut off the WI turbo pump and vacuum rack. Based on the probe data, the WI turbo pump almost certainly generates the small peak we see left at 45 Hz and the WI vacuum rack probably generates one of these leftover peaks in the 42-44 Hz range.

When we shut off the vacuum racks, we did it sequentially, with 10 minute spaces, which allowed us to match the disappearing peaks to specific racks. See Figure 5. This spectrogram shows us that each vacuum rack generates one distinct peak that is heard by the EIB episensor. Unfortunately, our data for the spectrograms was not good enough to identify the 42.3 Hz peak we saw disappear in the SEBDCE07 spectrum. This is not a great loss though because it is such a small peak.

3.3.1 The Mystery of the 41.8 Hz Line

Before the shut off test, it seemed clear that the BS vacuum rack was the source of the 41.8 Hz peak, which is probably the strongest in the entire forest. The line did disappear when we shut off the BS vacuum rack, which indicates that the BS vacuum rack is indeed a source, but only in SEBDCE09 (the vertical EIB episensor signal).

Strangely enough, the line remained in SEBDCE07 and 08 (horizontal). This tells us that there is another source at the same frequency. We went back through the probe data to find what we had missed and discovered that laser lab rack B1 also generates a peak at 41.8 Hz. See Figure 8.

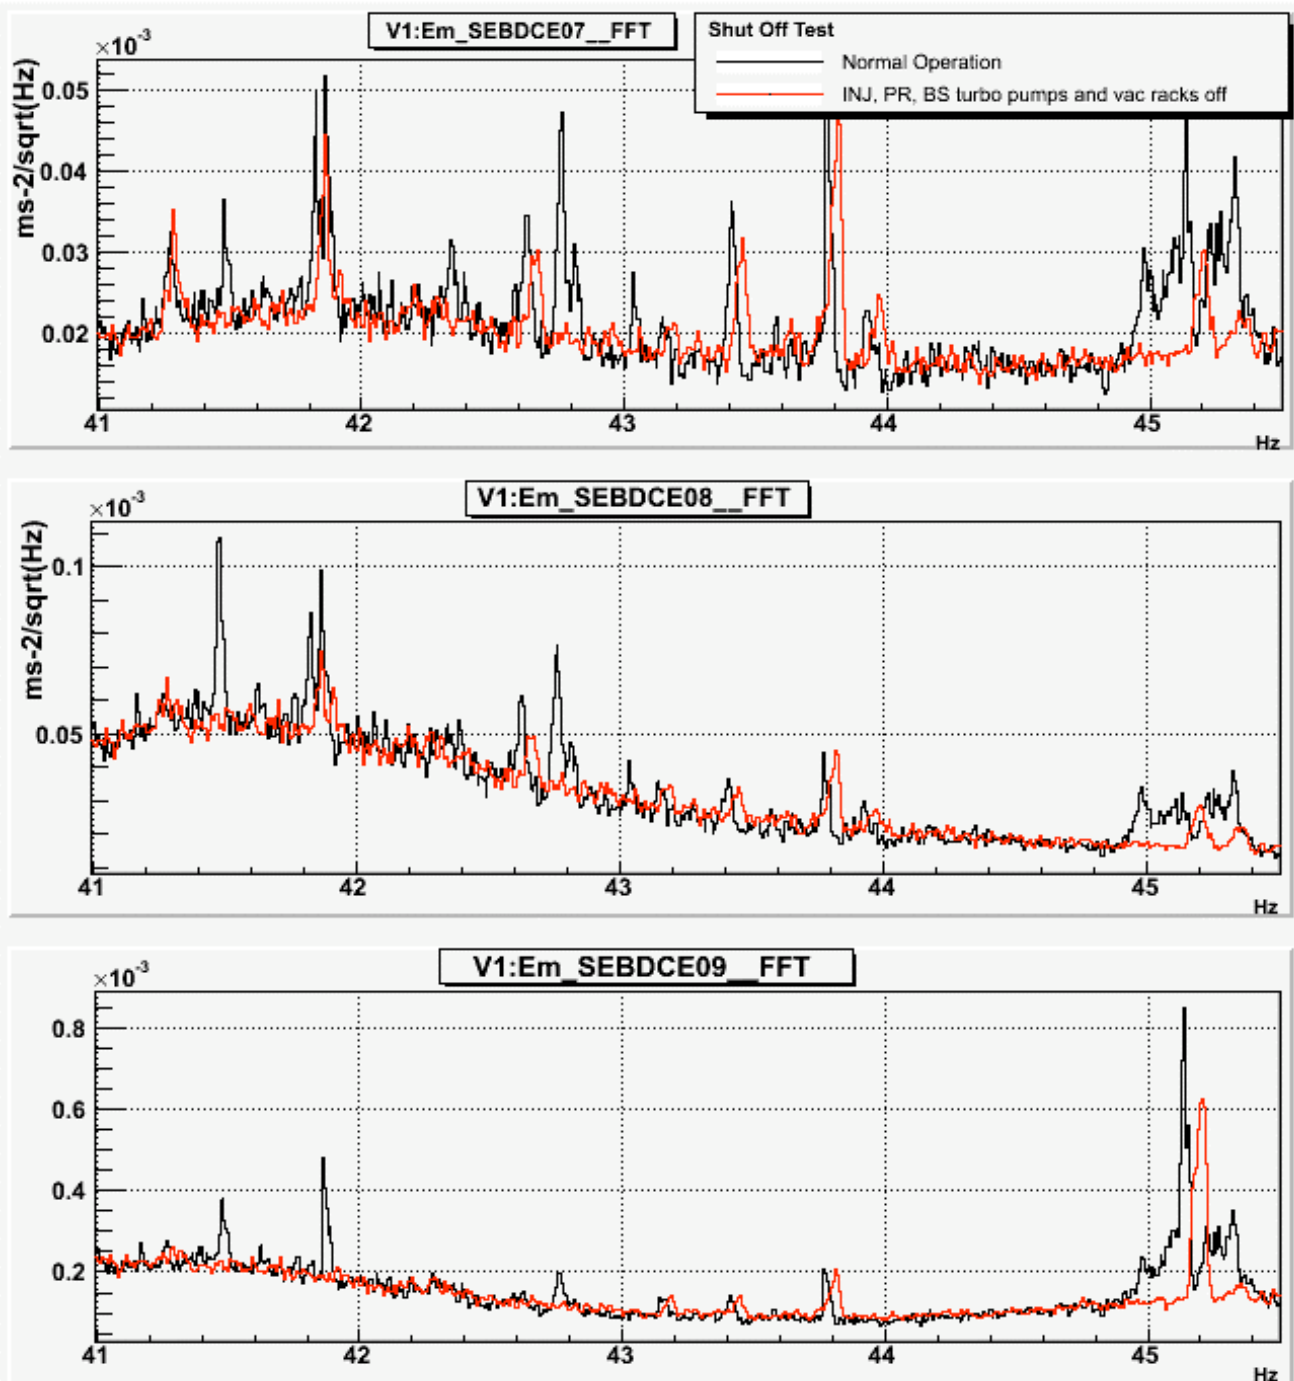
One possibility that would explain this is that these racks are not really as noisy as the probe makes them seem. Instead, the rack could be resonant at these peaks we think are fan noise. Because of the way we mounted the accelerometers (vertically, on top of the case), we should in fact pick up the resonances of the rack.

If it happens that one of these resonances is in this 40 - 45 Hz range, then it will show up as a strong peak on the probe, but there should be no coherence with the episensor. This is probably the cause of the large, incoherent peak we see at 41.5 Hz on laser lab rack B4.

3.3 Shut Off Tests

After about two weeks of probing, we were confident that the vacuum racks and the turbo pump cooling fans were the sources of almost all of the peaks in the forest, so we performed a shut off test. On July 3rd, we shut off the INJ, PR, and BS turbo pump cooling fans and the INJ, PR, and BS vacuum racks. See

In SEBDCE07 (East-West channel), we see two peaks: one smaller peak corresponding to the peak in the rack B1 probe channel and a larger peak which must be generated by the BS vacuum rack. We see strong coherence at the frequency of the rack B1 peak in both the horizontal and vertical EIB episensor channels, but the smaller peak generated by rack B1 is not really seen in SEBDCE09. This means that the rack B1 peak only couples in the horizontal directions and explains why we saw the line disappear only in the vertical episensor channel.



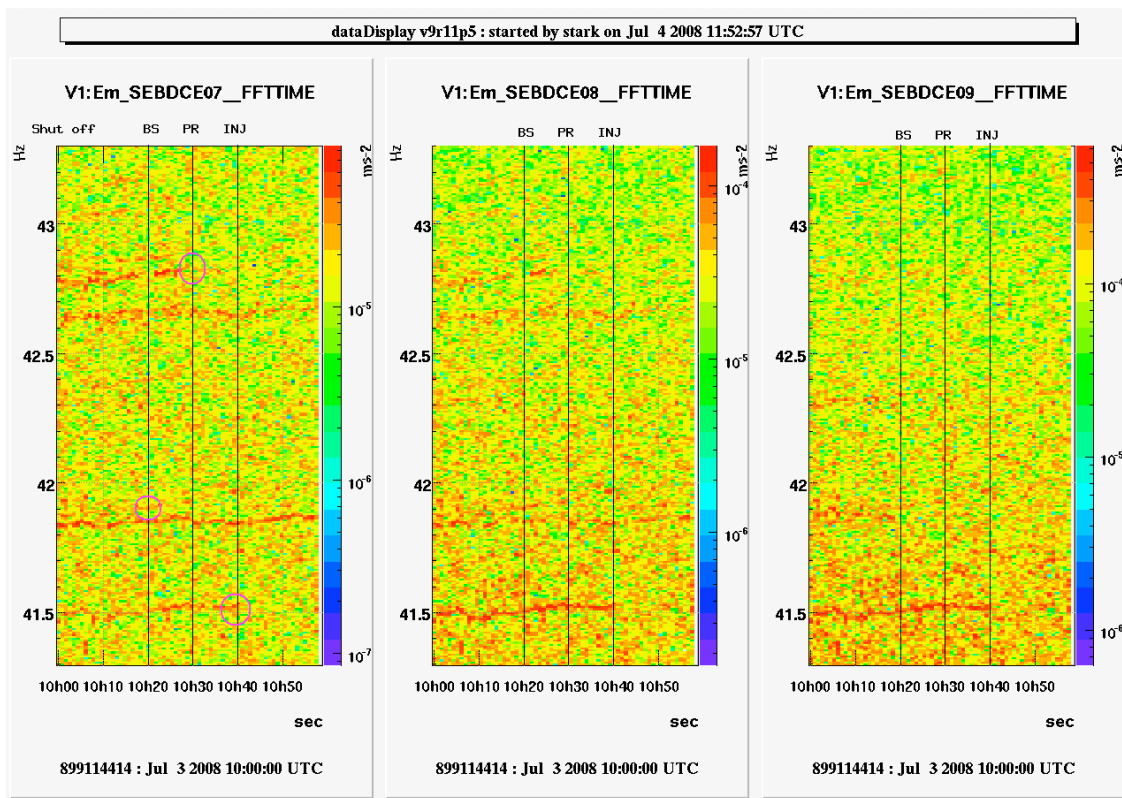


Figure 7: Spectrogram of the EIB episensor signals during the shut off test. Shows individual lines disappearing as racks are shut off separately.

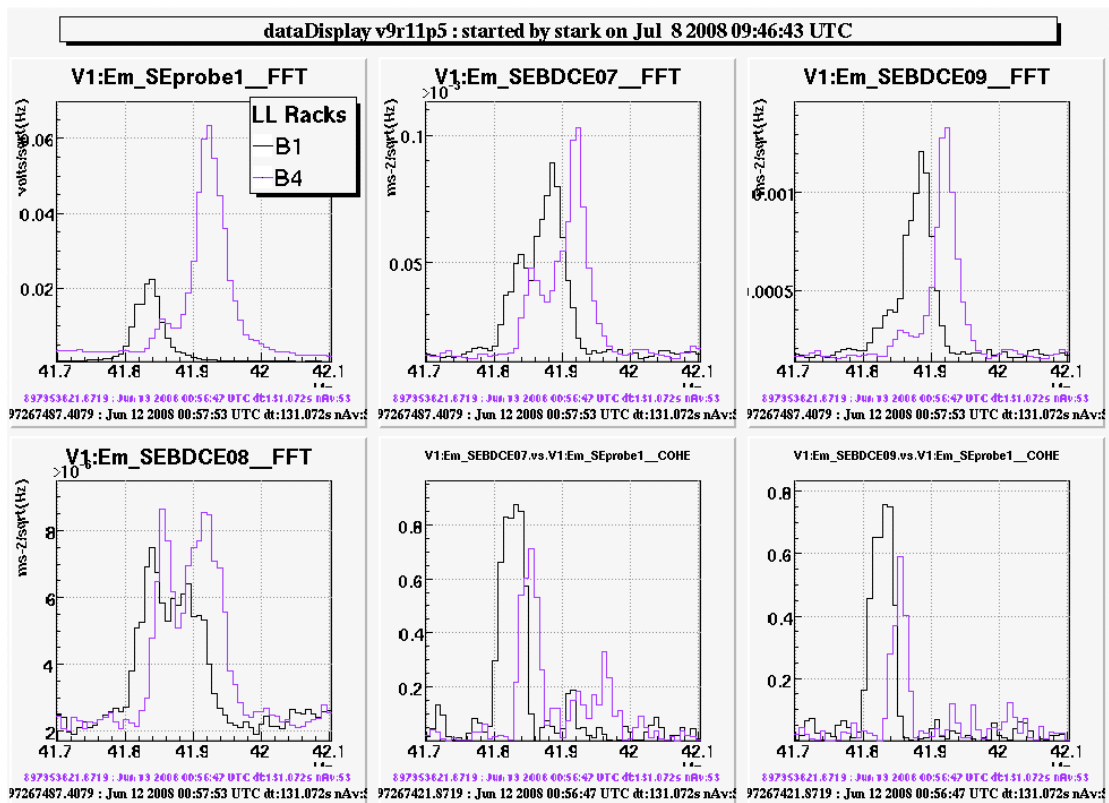


Figure 8: Laser lab rack probe data shows that rack B1 creates a peak at 41.8 Hz, but it only couples strongly in the horizontal directions.

3.3.2 Shut Off Test Results

Although the probe tests gave us a good idea of what we should look for, the shut off tests were the only way to positively identify sources. There was always a lot of guessing to do with the probe data, but the shut off gave us undeniable results.

With all of the data, we were able to identify five peaks in the forest. The results are shown in Table 1, along with the peaks we have not identified. Again it should be noted that one of these unknown peaks is most likely generated by the WI vacuum rack, but we were unable to determine which one.

Frequency (Hz)	Source
41.5	INJ vacuum rack
41.8	BS vacuum rack / LL rack B1
42.6	Unknown
42.8	PR vacuum rack
43.0	Unknown
43.4	Unknown
44.0	Unknown
45.0 / 45.2	Turbo pump cooling fans. Mostly INJ and PR.

3.3.3 Noise Coupling

Although we definitely identified five sources, this doesn't do us much good without knowing the coupling path. Otherwise, we can't be sure how to mitigate the noise.

Before we do anything, we have to figure out how the noise gets to the EIB (or interferometer, but we're already assuming this is how it enters the interferometer). The data we had did not tell us anything conclusive about the coupling path of the vacuum racks, but we were able to do something about the turbo pumps. See Figure 9.

In this figure, we have a clear view of the 45 Hz hump, in the EIB episensor and EIB accelerometer (SEIBLL02), the INJ tower accelerometer (SETOIN01), and the laser lab microphone (ACLALL01) spectra before (black) and during (purple) the shut off. In every sensor but the microphone, we see a distinct hump that disappears during the shut off. The signal remains unchanged in the microphone data, which tells us that the noise is not transmitted acoustically, but seismically.

3.3.4 Fan Harmonics

We also thought of considering harmonics of the fans to see if there are more lines at higher frequencies. This was mentioned in the second logbook entry by J. Marque [3], but the evidence was not conclusive then.

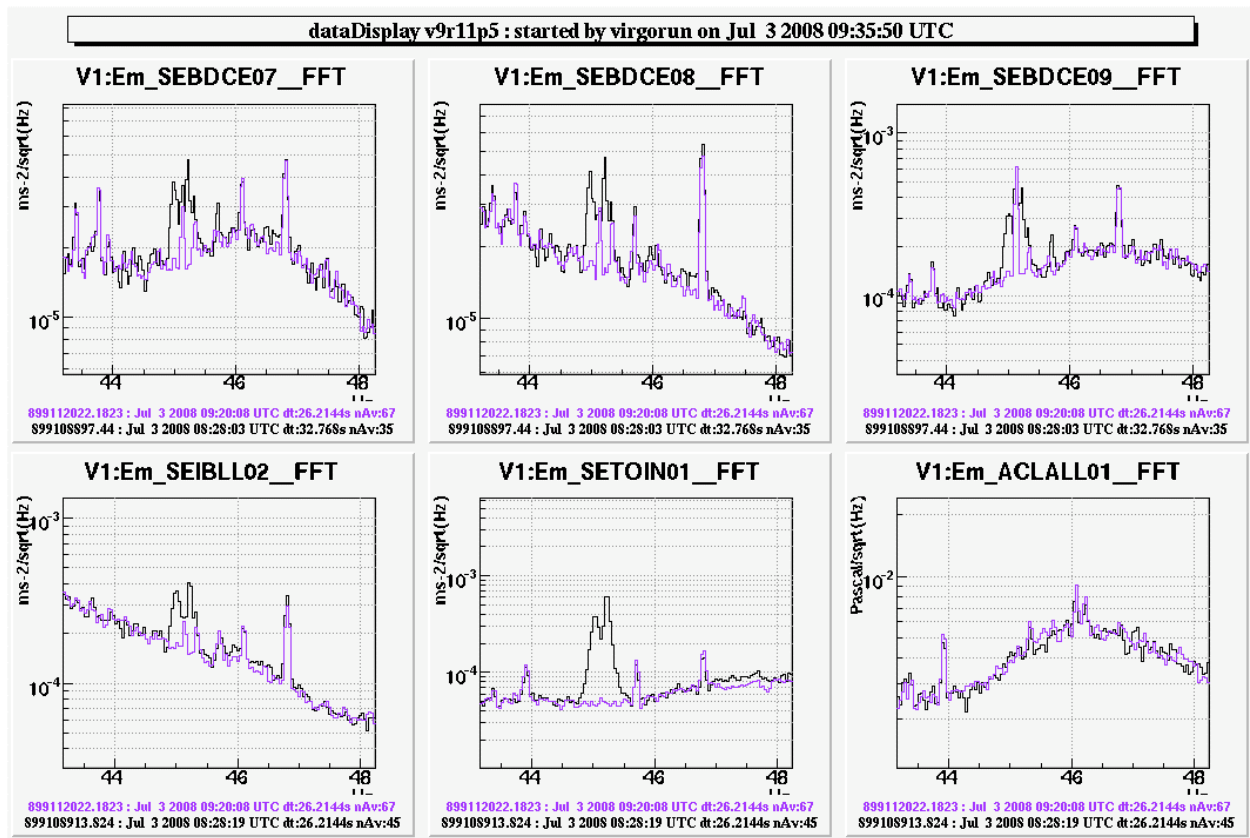


Figure 9: EIB episensor, INJ tower accelerometer, and LL microphone demonstrate that the turbo pump noise couples seismically.

We decided to look at the shut off test data, since a peak disappearing would prove that harmonics reach the EIB. We looked at the probe first and saw that the noise increased a lot when the rack was restarted. Because we see these harmonic peaks so strongly on the probe, we also checked the EIB episensor to see if they are transmitted. Zooming to 80 to 90 Hz to see the first harmonics of the fundamental fan frequencies, we saw at least one peak disappear, as seen in Figure 10. This peak corresponds to the first harmonic of the 41.5 Hz peak from the INJ vacuum rack.

This is very important because it tells us that some of the fans' harmonics are also transmitted to the EIB. We thought that this would not be the case since higher frequency seismic noise isn't transmitted as well, but it is obvious that at least the first harmonic of this fan noise reaches the EIB. Although this is a bad thing because it means more noise in the interferometer, it is also a good thing because it means that getting rid of the 41 - 45 Hz forest will reduce noise at higher frequencies as well.

3.4 Fan Tests

In each vacuum rack, there are two sets of four fans (middle and top), which means eight total in each rack. We decided to perform one last test to determine which set of fans is noisier. Ideally, we will be able to turn off all of the vacuum rack fans and completely remove these peaks from the dark fringe, but we don't know if that is possible yet. In the case that we can turn off only one set of fans, we want to know which set to turn off in order to reduce noise the most.

Using both the BS and PR vacuum racks, we turned off only the top fans, then only the middle fans,

and also both sets of fans. Using the probes that were attached to the racks during the test, we were able to see exactly which lines were generated by the top and the middle fans separately and the results also confirmed the results of the first shut off test. See Figure 10.

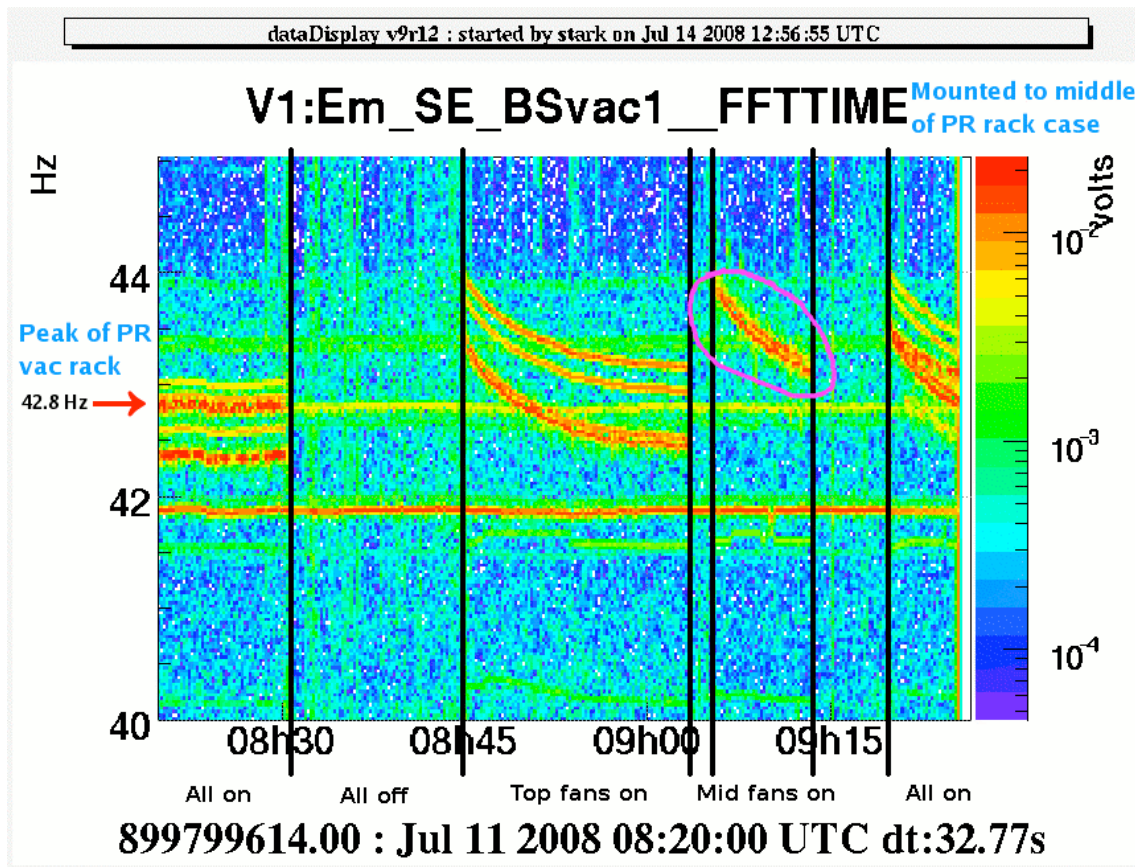


Figure 10: Spectrogram of BS vacuum rack probe shows that the primary peak is generated by the middle set of fans.

The one problem with the fan tests is that the fans run at a higher frequency when the rack is turned on, and then slow down to their resting frequency. Unfortunately, this slowing down process takes much longer than the time we allowed, and we had to extrapolate the lines to guess the final frequency. Luckily, this was not very difficult because the most prominent lines are so easy to identify. Figure 10 indicates where the noisiest line is on both racks and shows that they come from the middle set of fans on both the PR and BS vacuum racks.

3.5 Further Testing and Plans for Mitigation

We were able to identify several peaks in the forest, but much remains to be done in order to get rid of this noise.

The first task is to move the laser lab racks out of the laser lab. This has been planned for some time and should be finished before the interferometer goes online again. This action should reduce the 41.8 Hz peak and will hopefully reduce much of the background noise in the range of the EIB's resonance.

Another task is to perform temperature tests on the vacuum racks. The plan is to turn off the fans of each vacuum rack and monitor their temperature to make sure that they can survive without the active

cooling. We hope that we can turn off all of the fans in all vacuum racks permanently. At the very least, this will get rid of four peaks in the forest. If the vacuum rack fans need to be kept on, we should try to pad the racks to reduce any seismic coupling.

For the turbo fan cooling fans, we are working on a way to use them without rigidly connecting them to the tower. This way, the seismic coupling will be greatly reduced. Hopefully, we will be able to design a new mount for the fans that connects to the platform instead of to the tower.

Most importantly however, we want to improve the seismic isolation of the external injection bench. By at least modifying the resonance of the bench, we can greatly reduce this noise and improve sensitivity at low frequencies. The noise group at Virgo is currently investigating the properties of the optical benches to see if they can increase the resonant frequencies and isolate the bench any better.

4 Line Monitor

Last summer, a student from the same program built the Virgo Line Monitor Web Interface. This php noise analysis tool is very useful for quickly checking long term noise trends.

The Line Monitor (LM) works by connecting to a database that is updated in real time when the interferometer is on. A dedicated server stores rows of data in this database based on a trigger that looks for peaks. When this program finds a peak, it saves the frequency, GPS time, amplitude, and sound-to-noise ratio of the noise. This database allows the LM to work much faster by not having to sort through huge sets of Virgo data.

To use the LM, you select data by specifying a start time in GPS time and a duration in seconds for the time window, a frequency range, and a lower limit for the SNR. Then a plot is generated using this data and the user can change the axes to any of the four data fields: freq, gpstime, amp, and SNR. The default axes are freq vs. gpstime, because this shows noise trends over time.

The latest version of the LM needs an update. Several requests have been made from various members of the noise group. One problem is that the LM only allows input of GPS time, but this is very awkward, and users want to be able to input the time window in UTC time format. Another feature users would like is the ability to either specify a stop time or a duration (in seconds) for the time window.

The biggest problem with LM though is the database. It is currently handling one table that is almost an entire gigabyte of integers. It seems that no one ever cared to optimize this database because the variable types and indexes are badly chosen and have not been updated in a long time. As a result, the database has slowed down a lot as it has grown.

Unfortunately, I only had about a week and a half to work on this whole project because of other work so I didn't have time to fix up the database. I did, however, get to build a new version of the interface. The old code was good from a conceptual standpoint, but it was also very bulky, so I tried to clean that up. I also fixed little bugs from the previous version like making sure that the rest of the page is printed before the plot is processed.

A lot is left to be done on this project. The database should be optimized, re-indexed, and the main table should be partitioned by gpstime (say, into 3 month segments?). The interface is not finished

either. We should add some better validation rules to make sure users can't enter bad data, fix a few minor formatting issues, and possibly use Ajax to create a better display for the lock time periods.

References

- [1] The 41-44Hz forest, **J. Marque**; Virgo Logbook #20454.
- [2] Projection of EIB noise including also direct coupling, **I. Fiori**; Virgo Logbook #20548.
- [3] The 41-44Hz forest (2), **J. Marque**; Virgo Logbook #20457.