LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note **LIGO-T1500536-v3**

Following GW150914 Through the DARM Signal Chain

Matthew Evans

This is an internal working note of the LIGO project.

California Institute of Technology
LIGO Project, MS 18-34
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW17-161
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

WWW: http://www.ligo.caltech.edu/

Contents

1 Introduction **3**

2 Rogue Injections **4**

3 DARM Loop Signal Reconstruction **5**

3.1 DARM with Hardware Injection via the Electrostatic Drive (ESD) 10

3.2 DARM with Injection via the Photon Calibrator (PCAL) 14

4 Conclusions **16**

1 Introduction

An interesting signal, dubbed GW150914, appeared in the Advanced LIGO data in the early hours of September 14, 2015. This signal appears to be of astrophysical origin, possibly originating in the merger of a binary black hole system several hundred million parsecs from the Earth. For reasons which will become clear shorty, this potentially momentous signal was initially greeted with skepticism; many, if not all, of the scientist aware of it assumed that it was a fake, placed in the data to test the analysis pipelines. This initial reaction was natural because these fakes (known as "injections") appear regularly, while a real astrophysical signal has never been observed.

Injections come in a variety of flavors. The kinds of injections of interest here are those known as "hardware" injections, their close cousins the "blind" injections, and a previous unknown type now referred to as "rogue" injections. Hardware injections are signals added to the digital control system of a LIGO interferometer (the part of interest is known as the "DARM loop"), to test the response of the system and the analysis software. A careful public record is kept of these injected signals such that they will never be mistaken for a real signal.

Blind injections are performed in a manner identical to hardware injections, but the record of their presence is not public. They are intended to test not only the hardware and software infrastructure, but also the human parts of the detection process. Clearly, to have a good test, the scientists involved cannot know that the signal they are investigating is a fake. It should be noted, however, that "blind" injections are only blind *by convention*; evidence of their presence is available to any LSC member who dares to break the rules of the game.

Because GW150914 happened just before the O1 observing run was set to start, and the software required for performing hardware injections (blind and not) was still under development, an immediate check was made to ensure that the observed signal did not originate from an unintended injection (see EVNT log 11195). No trace of an injection was found.

Given that the standard injection methods had been ruled out, only a few other possibilities remained:

1. **G-Wave**: This is a gravitational wave signal
2. **Environmental**: The signal was caused by an external event simultaneously at both observatories
3. **Hardware Malfunction/Hack**: This signal was produced simultaneously in the hardware of both interferometers
4. **rogue Injection**: This signal was injected, but not in a standard way
5. **Frame Spoof**: The signal was added to the data after it left the interferometer control system

The remainder of this document will address only option 4, rogue injections. The reader may rest assured, however, that on similarly close inspection both options 2, 3 and 5 appear much less likely than option 1 (see EVNT log 11380, EVNT log 11376, EVNT log 11383, and T1500514).

On option 5; while not directly aimed at ruling out this option, the analysis presented here makes it clear that Frame Spoofing is a monumental task. Even if we ignore the substantial challenge of breaking into all of the relevant computer systems and finding all of the redundant data files (see EVNT log 11383), the 30+ channels which contain different version of DARM, analyzed herein, make this a daunting hack.

2 Rogue Injections

For the sake of this work, let us define a "rogue injection" as the addition of a signal which could mimic an astrophysical source to any point in the interferometer's digital control system. Keep in mind that this injection could happen as a result of malicious intent, or simply due to some unfortunate software bug.

The most obvious means of injecting a signal is through an "excitation point" in the control system, which is where hardware injections and blind injections come in, but no unintended excitation points were in use at the time of GW150914 (see EVNT log 11253). This does not, however, rule out a wide variety of bugs/hacks which could add a signal without using an excitation channel.

The vast majority of signals in the digital system deal with auxiliary controls, such as the suspension or seismic isolation systems, and there is no reason to expect a gravitational wave signal to appear there. Furthermore, these signals couple poorly to the main gravitational wave signal chain (know as DARM) by design, so a large disturbance would generally be required to cause a noticeable signal in DARM. For all such signals, a relatively simple search which looks for transient behavior at the time of GW150914 is sufficient to rule them out as the recipient of a rogue injection. This search has been performed (repeatedly), and no interesting transients were found (see EVNT log 11288, EVNT log 11267, and EVNT log 11414).

This leaves only the channels which naturally contain the gravitational wave (GW) signal as potential recipients of a rogue injection. These channels are part of the DARM control loop, and the presence of a GW signal at any point in the signal chain can be traced to the analog-to-digital converters (ADCs) which read the "gravitational wave photodiodes" (known as PD A and PD B). After entering the digital control system, the DARM signal passes through a number of time-domain filters before being sent to the actuators which keep the interferometer at its operating point via multiple digital-to-analog converters (DACs). A simple transient will naturally find the GW signal in these channels, so something more sophisticated is required to rule out rogue injections into the DARM loop.

In order to rule out rogue injections of all sorts, one need only ensure that the signals sent to the DACs is as expected given the signals which enter via the ADCs. That is, if the transfer functions from the ADCs which collect PD A and PD B to the n^{th} DAC signal are A_n and B_n, the DAC signal should be given by

$$DAC_n(t) = (A_n \otimes PD_A(t) + B_n \otimes PD_B(t)) \ . \tag{1}$$

Any rogue injection would cause the actual DAC signals to differ from the one computed according to this equation. As will be shown in more detail in the following sections, **there**

is no significant difference between the expected DAC signals and those recorded at the time of GW150914. This effectively rules out all kinds of injections into the digital control system, intentional or otherwise.

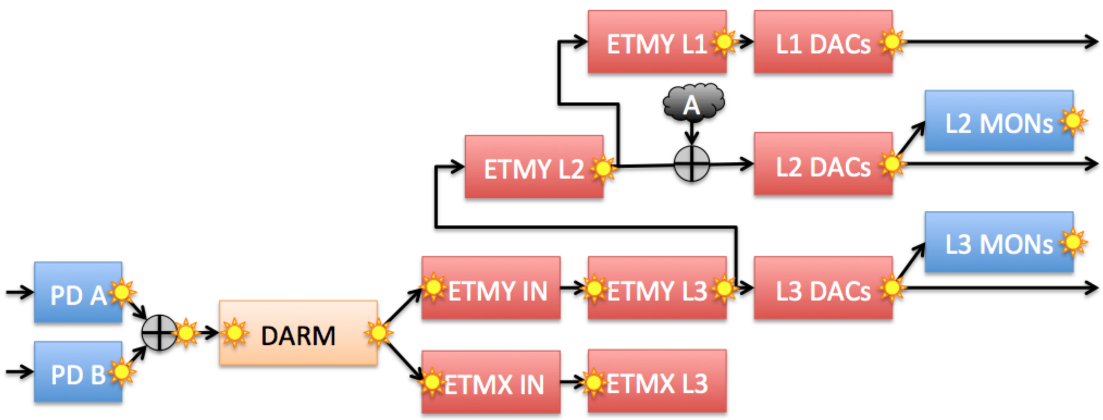

Figure 1: A sketch of the DARM signal path through the digital control system. The yellow markers indicate points where the signal is recorded. Each of the "DACs" and "MONs" markers represent 4 recorded signals (one for each quadrant), such that the total number of recorded DARM signals considered here is 32. Because of the alignment and damping signals which enter the L2 stage (marked by black cloud labeled "A"), the L2 reconstruction is imperfect, but it is sufficient to rule out any signal injection sufficient to cause GW150914.

3 DARM Loop Signal Reconstruction

While the end-to-end test suggested in equation 1 is sufficient to ensure that nothing was added to the signal chain, an even more stringent test is afforded by the fact that DARM is recorded at many points in the signal chain. The first of these is at the output of the filters which receive the PD A and PD B ADC signals (see figure 2).

The first record after the PDs is the SUM signal, which is followed by DARM IN (differing only by a gain factor, see figure 3). After the DARM filter, the signal is sent to the ETM suspensions according to the coefficients in the LSC output matrix (1 for ETMX and -1 for ETMY). The signal is recorded upon arrival in the suspensions, and as it appears at the input and output of the bottom stage of the suspension (i.e., the test mass, known as L3). Thus, there are 6 nearly identical copies of the signal which exits the DARM filter (aka DARM OUT). Only 2 of the 5 plots are shown, though all matched expectations (see figure 4).

The DARM signal which is sent to the L3 actuators is also sent up the suspension to the L2 and L1 stages (aka penultimate and upper-intermediate). Each stage performs some filtering, the result of which is recorded (see figure 5). There is a final stage of filtering before signals at each stage are sent to the DACs, generally used to compensate the analog dewhitening filters which follow Each of the signals sent to the DACs is recorded, with 4

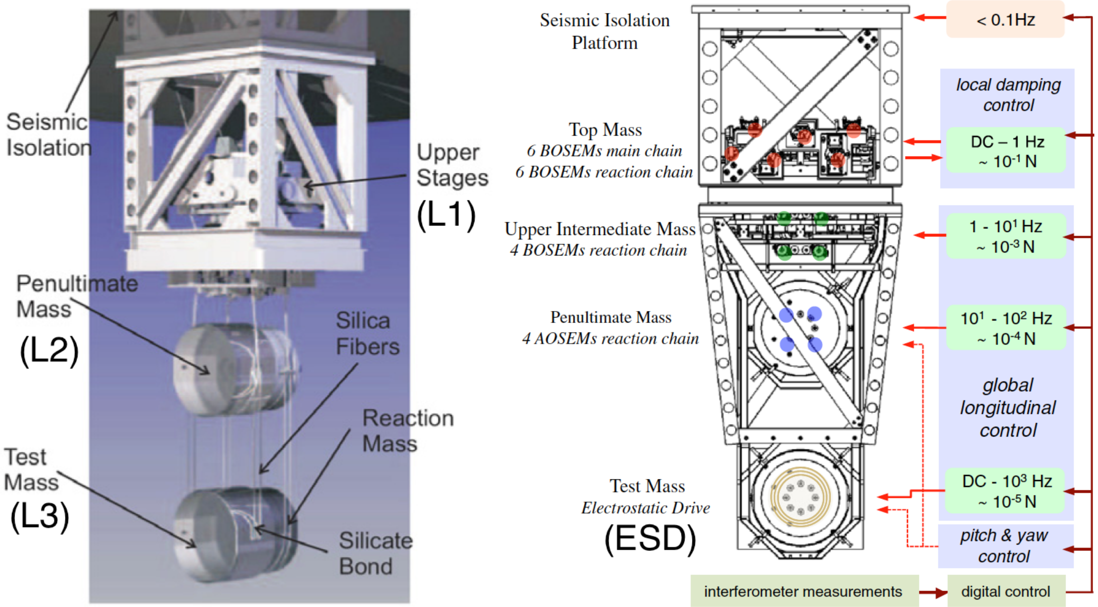

Figure 2: The Advanced LIGO test mass suspension, as shown in CQG 29 115005. Note that the electrostatic drive (ESD) is acts directly on the test mass, which is also known as "L3" in the sequential naming convention of the digital control system.

for each suspension stage due to the quadrant based construction of the suspensions. Two representative signals are shown in figure 6.

Finally, after being converted to analog by the DACs, all suspension channels have monitors which record the signal actually sent to the actuators (figure 7). The monitors are there to detect differences between what is recorded at the DAC, and what actually went to the suspension. As such, they also see anything added via the test inputs to the suspension electronics. These monitor signals are digitized and recorded, though unfortunately the L3 monitors were recorded at 256Hz and are of little use for detecting injections.

This analysis shows that for GW150914 all 30 of the signals which can be calculated from PD A and PD B are consistent with normal signal propagation (e.g., they show no sign of an injected signal).

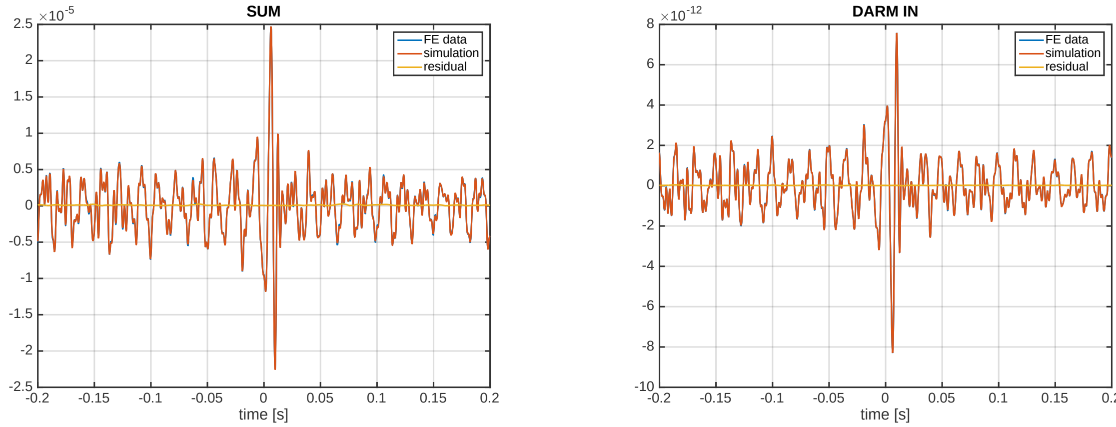

Figure 3: GW150914: The first 2 signals after the PDs (SUM and DARM IN). The blue curve (labeled "FE data") is the data recorded by the control system (CDS), while the orange curve (labeled "simulation") is the curve calculated from the PDs. They are essentially identical, as indicated by the residual, which is the difference between them. Note that "FE" is short for "front-end"; the name given to the real-time digital control computers.

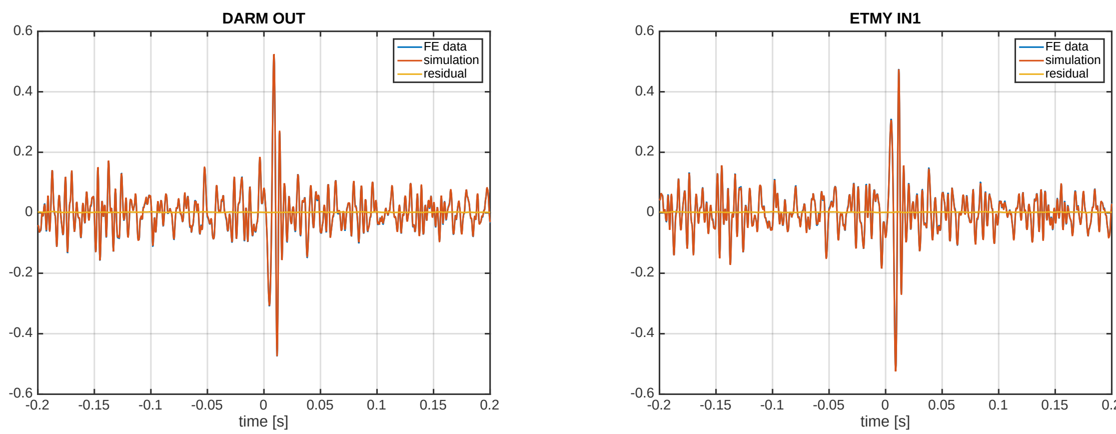

Figure 4: GW150914: The signal after the DARM filter, and the input to the ETMY suspension. Between the DARM OUT and ETMY IN1 points in the digital signal chain, the signal is multiplied by a constant matrix element (-1 in this case) and passed across the reflective memory network from the corner station to the Y-end.

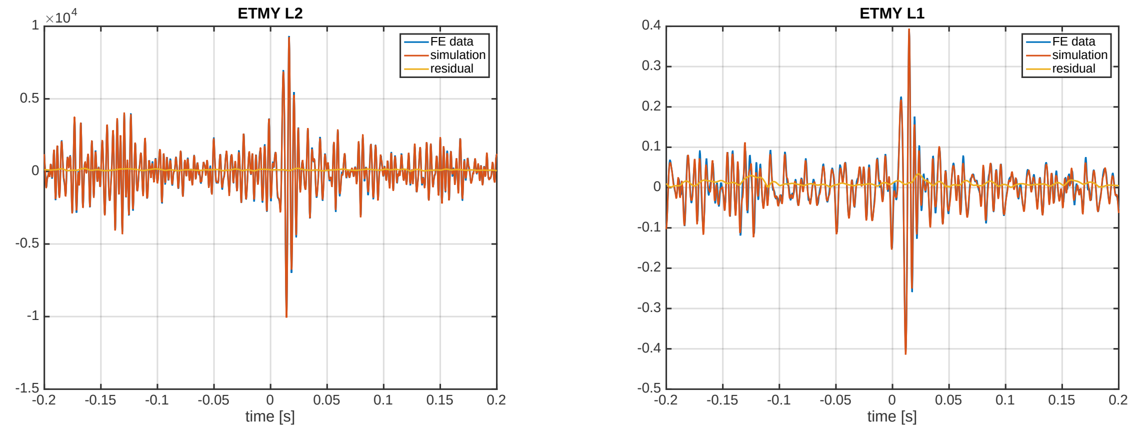

Figure 5: GW150914: The signal after the L2 and L1 filters as it moves up the ETMY suspension. In both plots the residual is visibly different from zero, mostly due to the sample rate conversions required to make the comparison. (All signals are converted to 16384Hz, the DARM sample rate, in order to compute the difference.) The **"residual"** is the **absolute value of the difference low-pass filtered at 100Hz**, giving it a smooth shape even where the signal changes quickly.

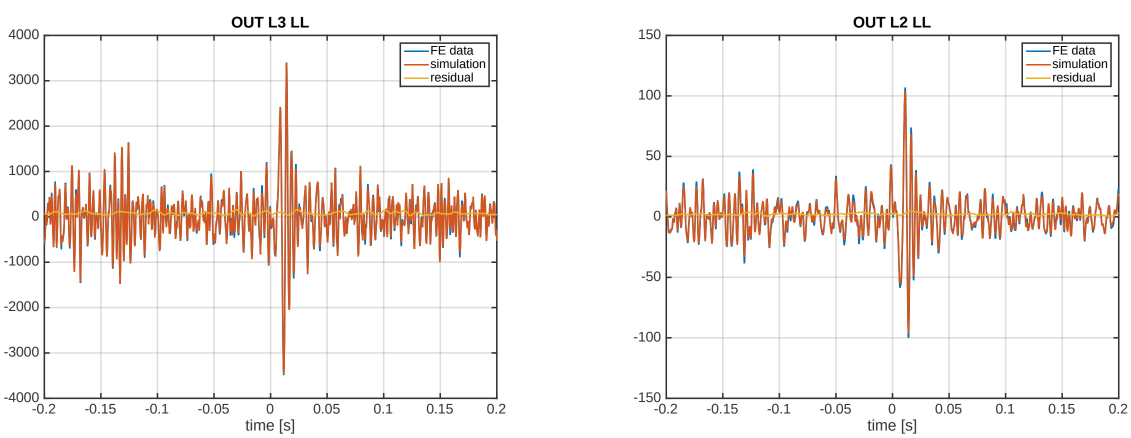

Figure 6: GW150914: The signals as they are send to the DACs. Only one quadrant for the L3 and L2 stages are shown. For L3 (the test mass, with its electro-static drive), all signals are identical. For L2, on the other hand, the actuators are also used for alignment and damping. The lower-left quadrant (shown here) suffers relatively minor contamination from these added signals.

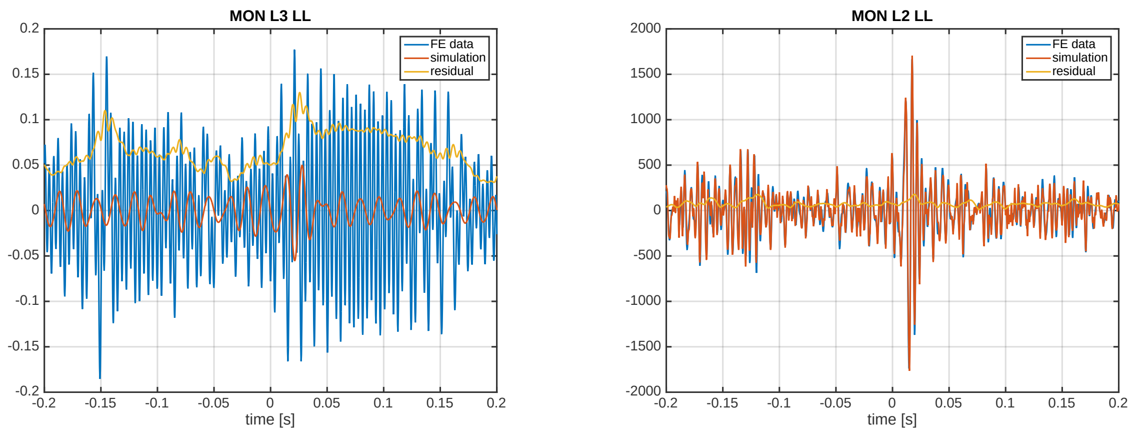

Figure 7: GW150914: The signals as seen by the actuation signal monitors. Unfortunately the L3 monitors (those of the ESD, at the test mass level) were recorded at 256Hz and are of little use for detecting injections. (Much of the blue FE signal for the L3 monitor is actually an up-sampling artifact, resulting from this low sample frequency.) For L2, on the other hand, the monitors clearly show the actuation above the noise.

3.1 DARM with Hardware Injection via the Electrostatic Drive (ESD)

To demonstrate the effectiveness of this technique, the same analysis is applied to a time when a hardware injection designed to mimic GW150914 was performed (GPS 1128125224). Figures 8 through 11 show the results.

The injection can clearly be seen in figures 9 through 11, since the injected signal is added between DARM OUT and the suspension inputs. This is to be expected since the DARM loop gain is less than unity above 40 Hz, such that the signal required to produce an event is inevitably larger than the response of the control system to that event. (On a similar note, Andy Lundgren noted in EVNT log 11335, that since the control system have several notch filers where the signal content is essentially zero, any injection is painfully obvious in that it has content at the notch frequencies. A similar analysis was also conducted by Shivaraj Kandhasamy in EVNT log 11336.)

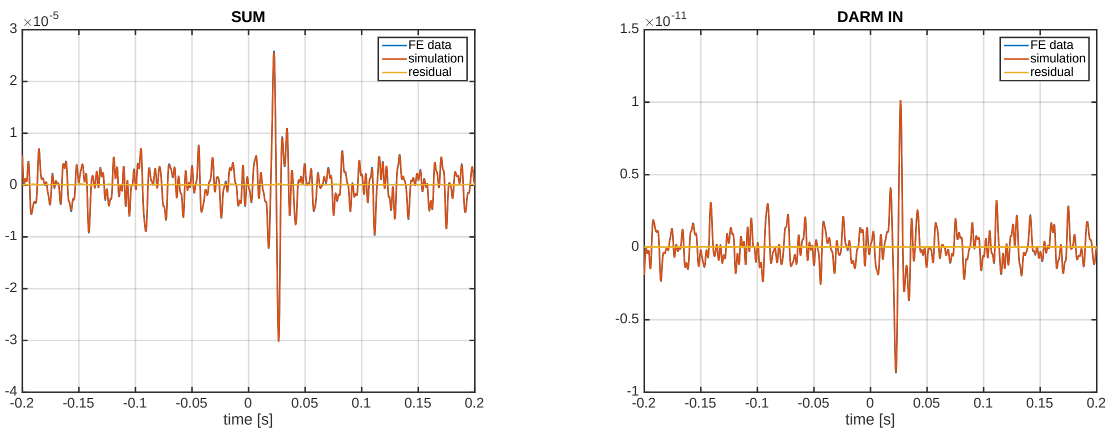

Figure 8: Hardware Injection: The first 2 signals after the PDs (SUM and DARM IN). The injection happens later in the chain, so these signal are as expected.

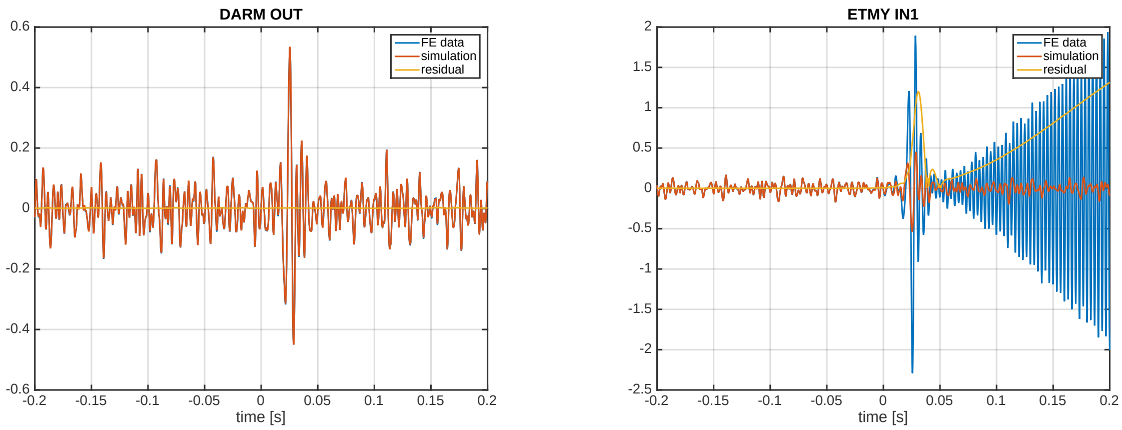

Figure 9: Hardware Injection: The signal after the DARM filter, and the input to the ETMY suspension. The hardware injection enters the signal chain just after the DARM filter (e.g., after DARM OUT, but before ETMY IN1). Note that the injection appears as a clear difference between the expected and recorded ETMY signal. (The long ringing transient present in the ETMY plot but not in later signals is due to the anti-notch filters required to cancel the notch filters which appear later in the suspension.)

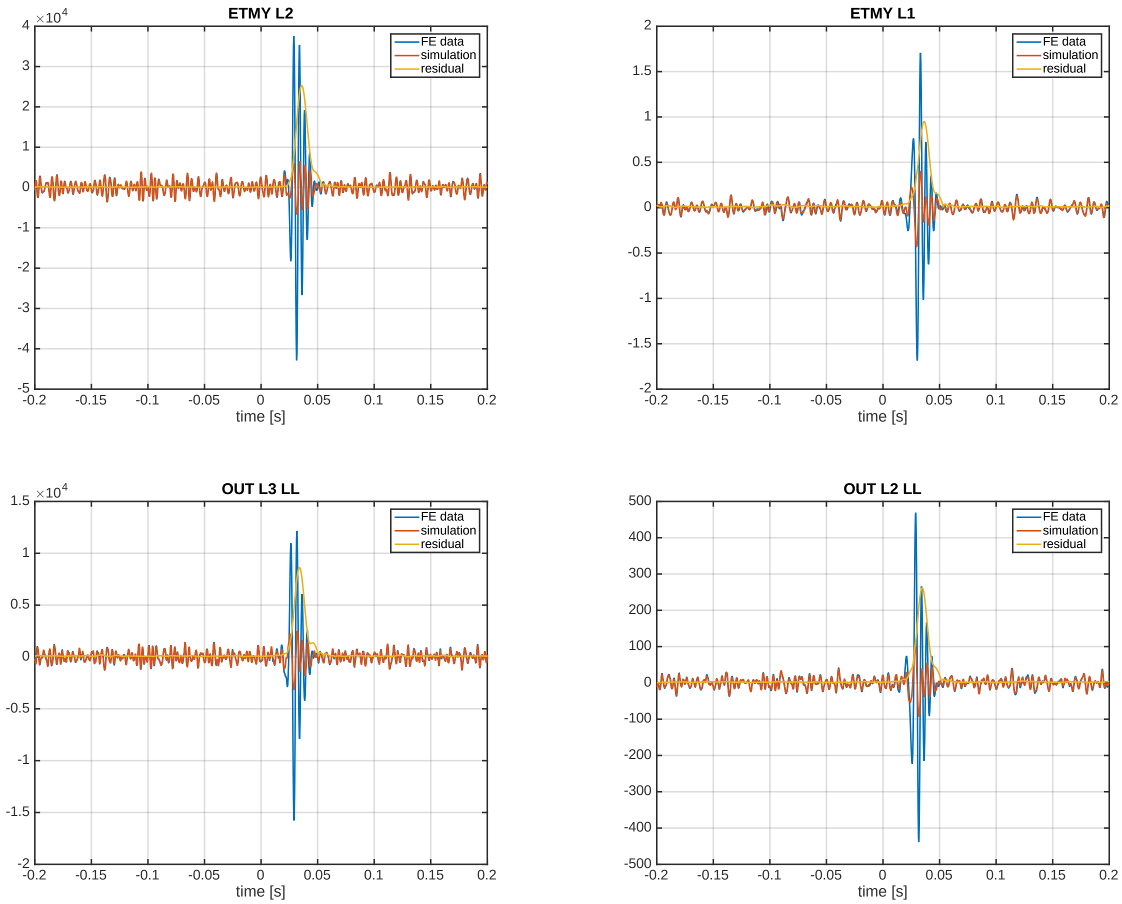

Figure 10: Hardware Injection: The signal after the L2 and L1 filters as it moves up the ETMY suspension and is sent to the DACs. The signal injected is not an exact match to GW150914, but has similar frequency content and amplitude (compare the orange traces here and in figure 9 with those of figures 4, 5, and 6). **The signal required to produce this hardware injection is easily seen in all channels.**

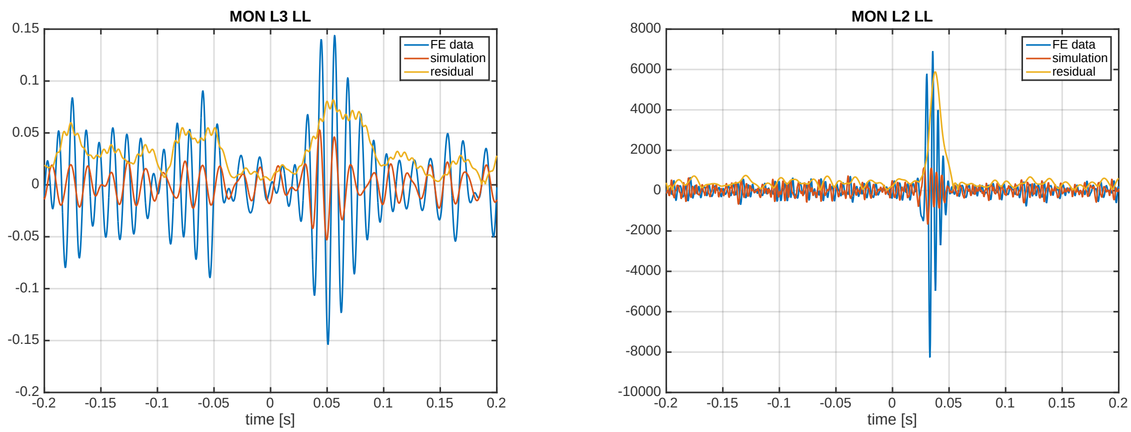

Figure 11: Hardware Injection: The signals as seen by the actuation signal monitors. Unfortunately the L3 monitors were recorded at 256Hz and are of little use for detecting injections. For L2, on the other hand, the monitors clearly show the actuation above the noise.

3.2 DARM with Injection via the Photon Calibrator (PCAL)

As a second demonstration, the same analysis is applied to a time when an injection designed to mimic GW150914 was performed via the photon calibrator (GPS 1128303099). This injection does not happen in the DARM loop, and thus more closely resembles a real gravitational wave event (since both essential apply a force to the interferometer optics). Figure 12 shows the results (somewhat abridged since they are all as expected).

Finally, figure 13 shows that this injection clearly appears in the PCAL monitor channel.

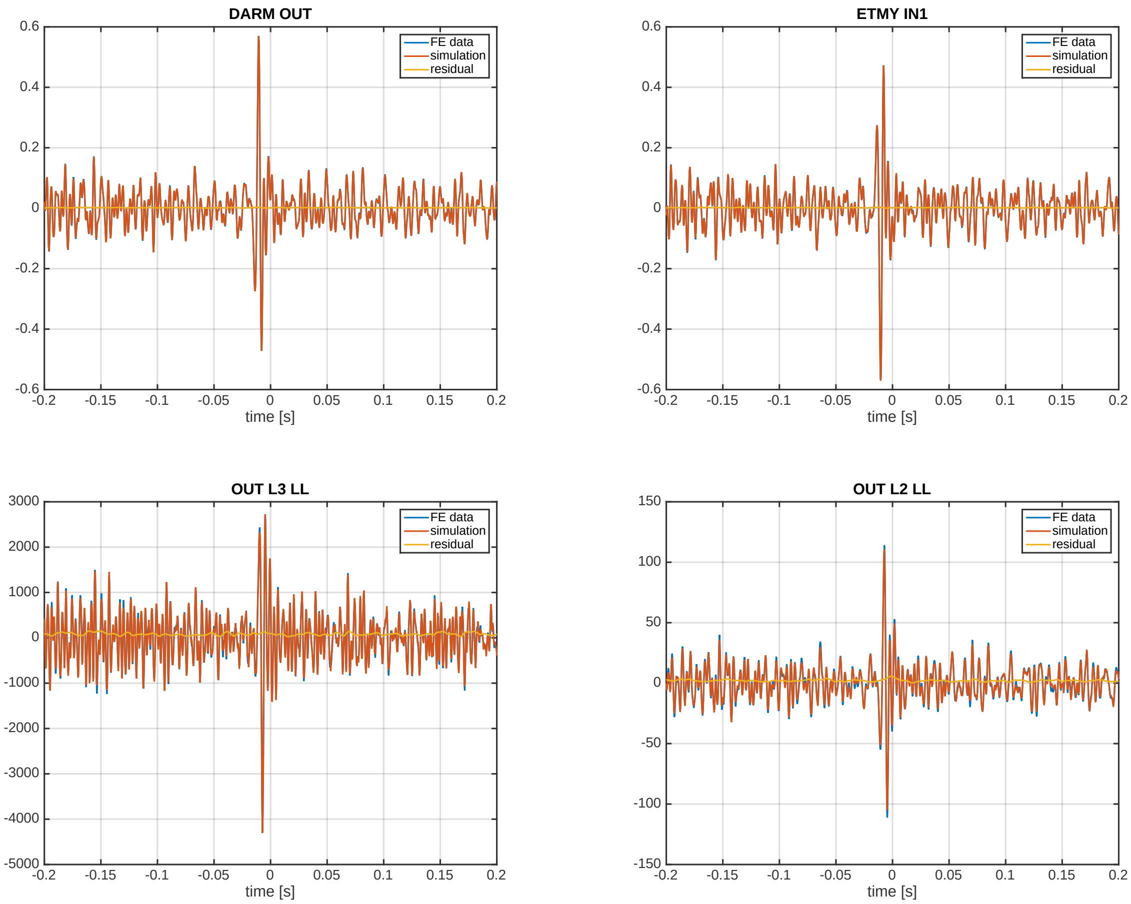

Figure 12: PCAL Injection: The signal after the DARM filter, at the input to the ETMY suspension, and at the L3 and L2 DACs.

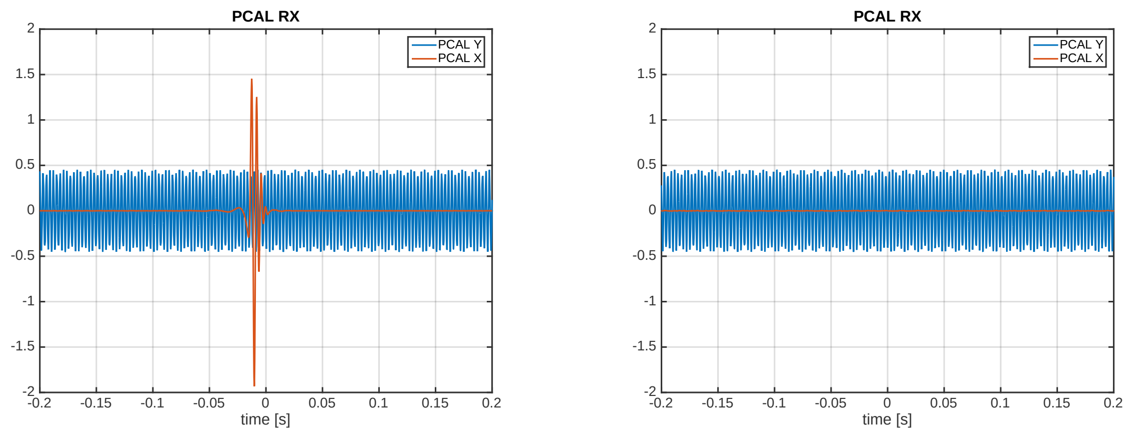

Figure 13: PCAL Injection: The output of the PCAL system is monitored and recorded. The plot on the left shows the usual calibration lines (blue) and the PCAL injection (orange). The plot on the right, showing data at the time of GW150914, shows only the calibration lines.

4 Conclusions

The possibility of GW150914 being an injection, sanctioned, blind, accidental, malicious or otherwise, was likely considered by all LSC scientists as quite likely when they first heard of the event. Given the fact that we had not yet started O1, and that the hardware injection software was being tested in the days around GW150914, this was a natural and prudent response.

While many checks have been done to ensure that GW150914 was not an intentional injection or the result of a software problem with the injection system, the analysis presented here takes an agnostic approach to the possible sources of an injection and simply **rules out any addition of a signal to the digital DARM signal chain** during GW150914.

To show how this result would be different in the case of an injection, results from a hardware injection after GW150914 are also shown. The presence of the injected signal, designed to mimic GW150419, is clearly evident.

A third set of plots shows an injection made with the photon calibrator. Since this injection is not in the digital DARM loop, but rather takes the form of a force acting on an interferometer optic, analysis of the DARM signal chain does not find an injection (which is instead evident in the photon calibrator monitor channels).

Finally, this analysis also makes it clear that adding a signal to the data after it is recorded is a monumental task (option 5 in the introduction, "Frame Spoof"). Since DARM is recorded in many channels, separated by inscrutably named matrix elements and digital filters with complicated transfer-functions, a proper fake would require intimate knowledge of the digital system's inner workings, as well as access to up-to-date configuration files and live parameter values (since the gains and filters change regularly during commissioning). It is doubtful that even a large conspiracy of LIGO insiders could manage a frame spoof.